Rotation Invariant Spherical Harmonic Representation of 3D Shape Descriptors

Michael Kazhdan, Thomas Funkhouser, and Szymon Rusinkiewicz

Department of Computer Science, Princeton University, Princeton NJ

Abstract

One of the challenges in 3D shape matching arises from the fact that in many applications, models should be considered to be the same if they differ by a rotation. Consequently, when comparing two models, a similarity metric implicitly provides the measure of similarity at the optimal alignment. Explicitly solving for the optimal alignment is usually impractical. So, two general methods have been proposed for addressing this issue: (1) Every model is represented using rotation invariant descriptors. (2) Every model is described by a rotation dependent descriptor that is aligned into a canonical coordinate system defined by the model. In this paper, we discuss the limitations of canonical alignment and present a new mathematical tool, based on spherical harmonics, for obtaining rotation invariant representations. We describe the properties of this tool and show how it can be applied to a number of existing, orientation dependent, descriptors to improve their matching performance. The advantage of this is twofold: First, it improves the matching performance of many descriptors. Second, it reduces the dimensionality of the descriptor, providing a more compact representation, which in turn makes comparing two models more efficient.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Methodology and Techniques

1. Introduction

Over the last decade, tools for acquiring and visualizing 3D models have become integral components of data processing in a number of disciplines, including medicine, chemistry, architecture and entertainment. With the proliferation of these tools, we have also witnessed an explosion in the number of available 3D models. As a result, the need for the ability to retrieve models from large databases has gained prominence and a key concern of shape analysis has shifted to the design of efficient and robust matching algorithms.

One of the principal challenges faced in the area of shape matching is that in many applications, a model and its image under a similarity transformation are considered to be the same. Thus, the challenge in comparing two shapes is to find the best measure of similarity over the space of all transformation. The need for efficient retrieval makes it impractical to explicitly query against all the transformations, and two different solutions have been proposed:

- **Normalization**: Shapes are placed into a canonical coordinate frame (normalizing for translation, scale and rotation) and two shapes are assumed to be near-optimally aligned when each is in its own frame. Thus, the best measure of similarity can be found without explicitly trying all possible transformations.

- **Invariance**: Shapes are described in a transformation invariant manner, so that any transformation of a shape will be described in the same way, and the best measure of similarity is obtained at *any* transformation.

We have found that while traditional methods for translation and scale normalization provide good matching results, methods for rotation normalization are less robust and hamper the performance of many descriptors.

In this paper we present a novel tool, called the *Spherical Harmonic Representation*, that transforms rotation dependent shape descriptors into rotation independent ones. This tool contributes to the challenges of designing effective shape retrieval algorithms in three ways. First, it is a general tool that can be applied to many existing shape descriptors. Second, for most shape descriptors, the spherical har-

monic representation provides better matching results than those obtained by rotation normalization. Finally, the spherical harmonic representation provides a reduction in the dimensionality of the shape descriptor, thereby reducing both the space for storage and the time for comparison – key properties for the implementation of interactive shape retrieval systems.

The rest of this paper is structured as follows: In Section 2 we describe previous work in the area of shape retrieval. The spherical harmonic representation is presented in Section 3, which reviews the principal properties of spherical harmonics and provides a method for obtaining rotation invariant representations of spherical-based shape descriptors. In Section 4 we describe the mathematical properties of the spherical harmonic representation and discuss questions of invertibility. We provide a generalization of our method to voxel grids in Section 5. In Section 6 we provide empirical results comparing matching results of normalized descriptors with their rotation invariant representations. We provide an analysis of these results in Section 7 and we conclude in Section 8 by summarizing our results and discussing topics for future work.

2. Related Work

The problem of shape matching has been well studied in the graphics/vision literature and many methods for evaluating model similarity have been proposed. This paper is motivated by the increased availability and accessibility of 3D models, and focuses on the problem of shape retrieval from within large databases of models. In this context, the challenge is to provide a robust and efficient method for computing model similarity.

To address this challenge, many methods have focused on separating the matching problem into two components: An offline step, in which abstracted distinguishing information is extracted from each model independently, and an online step, in which the information between two models is compared. In order to allow for efficient retrieval, the offline step is usually designed to extract information which allows for simple and efficient comparison between models. In particular, many existing methods describe a 3D shape with an abstracted *shape descriptor* that is represented as a function defined on a canonical domain. Shapes are then compared by computing the difference between their descriptors, so that no explicit establishing of correspondences is necessary, and the online process can be efficient.

However, in the context of shape retrieval, one of the principal difficulties faced by these approaches is that a model and its image under a similarity transformation are considered to be the same. Thus, the challenge in comparing two models is to find the best measure of similarity over the space of all transformations. This challenge has been addressed in two different ways:

- Normalizing the models by finding a canonical transformation for each one.
- Characterizing models with a transformation invariant descriptor so that all transformations of a model result in the same descriptor.

(While explicitly solving for the optimal transformation using either exhaustive search or methods such ICP [13, 14], the Generalized Hough Transform[15], or Geometric Hashing [16], are also possible, these approaches cannot be applied to database retrieval tasks since the online comparison of models becomes inefficient.) Many hybrid methods exists and a few representative examples are shown in Table 1, which describes how these methods address translation, scale and rotation.

Representation	Tr	Sc	Rot
Crease Histograms [2]	I	N	I
Shape Distributions [3]	I	N	I
Extend Gaussian Images [4]	I	N	N
Shape Histograms [5] (Shells)	N	N	I
Shape Histograms [5]	N	N	N
Spherical Extent Functions [6]	N	N	N
Wavelets [7]	N	N	N
Reflective Symmetry Descriptors [8]	N	N	N
Higher Order Moments [9]	N	N	N
Exponentation EDT [12]	N	N	N

Table 1: *A summary of a number of shape descriptors, showing if they are (N)ormalized or (I)nvariant to each of translation, scale and rotation.*

In general, models are normalized by using the center of mass for translation, the root of the average square radius for scale, and principal axes for rotation. We have found that while the methods for translation and scale normalization are robust for whole object matching [10], rotation normalization via PCA-alignment does not provide provide a robust normalization for many matching applications. This is due to the fact that PCA-alignment is performed by solving for the eigen-values of the covariance matrix. This matrix captures only second order model information, and the assumption in using PCA is that the alignment of higher frequency information is strongly correlated with the alignment of the second order components. (Appendix A provides an analysis of this from a signal processing framework.) We have found that for many shape descriptors this assumption does not hold, and the use of principal axes for alignment hampers the performance of these descriptors.

Many of the descriptors that have used PCA-alignment represent a 3D shape as either a spherical function or a voxel grid, which rotates with the model. Examples of such descriptors have included:

- The Extended Gaussian Image [4], which describes the distribution of normals across the surface of the model
- Shape Histograms [5], which describe the distribution of points on the model across all rays from the origin
- Spherical Extent Functions [6], which describe the maximal extent of a shape across all rays from the origin
- Reflective Symmetry Descriptors [8], which describe the reflective self-similarity of a shape with respect to reflections about all planes through the origin
- The voxel description of Funkhouser et al. [12], which describes a model by computing the negative exponential of its Euclidean Distance Transform

For these type of descriptors, we propose a solution to the rotation problem by providing a mathematical tool, based on spherical harmonics, for obtaining a rotation invariant representation of the descriptors. Our approach is a generalization of the Fourier Descriptor [11] method to the sphere, characterizing spherical functions by the energies contained at different frequencies. This idea was intially proposed in [12], and this paper presents a detailed description of the descriptor, its properties, and empirical results demonstrating its efficacy in improving the matching performance of a number of existing shape descriptors.

3. Spherical Rotation Invariance

In this paper, we present a tool for transforming rotation dependent spherical and voxel shape descriptors into rotation invariant ones. The key idea of our approach is to describe a spherical function in terms of the amount of energy it contains at different frequencies. Since these values do not change when the function is rotated, the resulting descriptor is rotation invariant. This approach can be viewed as a generalization of the Fourier Descriptor method [11] to the case of spherical functions.

3.1. Spherical Harmonics

In order to be able to represent a function on a sphere in a rotation invariant manner, we utilize the mathematical notion of spherical harmonics to describe the way that rotations act on a spherical function. The theory of spherical harmonics says that any spherical function $f(\theta, \phi)$ can be decomposed as the sum of its harmonics:

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} a_{lm} Y_l^m(\theta, \phi).$$

(This decomposition is visualized in step (1) of Figure 1.) The key property of this decomposition is that if we restrict to some frequency l, and define the subspace of functions:

$$V_l = \text{Span}(Y_l^{-l}, Y_l^{-l+1}, \ldots, Y_l^{l-1}, Y_l^{-l})$$

then:

- **V_l is a Representation For the Rotation Group**: For any function $f \in V_l$ and any rotation R, we have $R(f) \in V_l$.
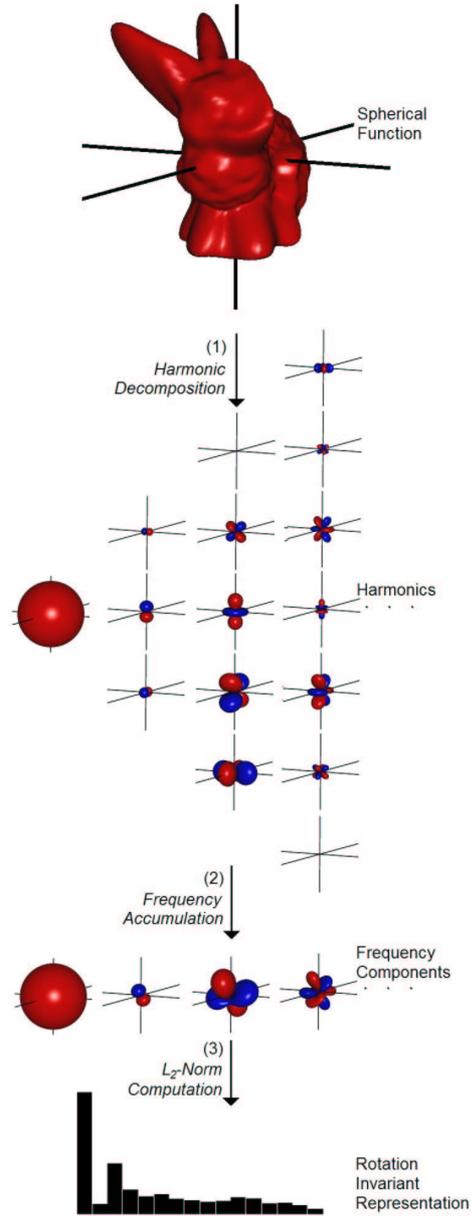
Figure 1: *We compute a rotation invariant descriptor of a spherical function by (1) decomposing the function into its harmonics, (2) summing the harmonics within each frequency, and (3) computing the norm of each frequency component. (Spherical functions are visualized by scaling points on the unit sphere in proportion to the value of the function at that point, where points with positive value are drawn in light gray and points with negative value are drawn in dark gray.)*

This can also be expressed in the following manner: if π_l is the projection onto the subspace V_l then π_l commutes with rotations:

$$\pi_l(R(f)) = R(\pi_l(f)).$$

- **V_l is Irreducible**: V_l cannot be further decomposed as the direct sum $V_l = V'_l \oplus V''_l$ where V'_l and V''_l are also (nontrivial) representations of the rotation group.

The first property presents a way for decomposing spherical functions into rotation invariant components, while the second property guarantees that, in a linear sense, this decomposition is optimal.

3.2. Rotation Invariant Descriptors

Using the properties of spherical harmonics, and the observation that rotating a spherical function does not change its L_2-norm we represent the energies of a spherical function $f(\theta, \phi)$ as:

$$\text{SH}(f) = \{\|f_0(\theta, \phi)\|, \|f_1(\theta, \phi)\|, \ldots\}$$

where the f_l are the frequency components of f:

$$f_l(\theta, \phi) = \pi_l(f) = \sum_{m=-l}^{m=l} a_{lm} Y_l^m(\theta, \phi)$$

(shown in steps (2) and (3) of Figure 1.)

This representation has the property that it is independent of the orientation of the spherical function. To see this we let R be any rotation and we have:

$$\begin{aligned}
\text{SH}(R(f)) &= \{\|\pi_0(R(f))\|, \|\pi_1(R(f))\|, \ldots\} \\
&= \{\|R(\pi_0(f))\|, \|R(\pi_1(f))\|, \ldots\} \\
&= \{\|\pi_0(f)\|, \|\pi_1(f)\|, \ldots\} = \text{SH}(f)
\end{aligned}$$

so that applying a rotation to a spherical function f does not change its energy representation.

3.3. Further Quadratic Invariance

We can make our representation still more discriminating by refining the case of the second order component. Using the results from Appendix A we know that the L_2-difference between the quadratic components of two spherical functions is minimized when the the two functions are aligned with their principal axes. Thus, instead of describing the constant and quadratic components by the two scalars $\|f_0\|$ and $\|f_2\|$, we can represent them by the three scalars a_1, a_2, and a_3, where after alignment to principal axes:

$$f_0 + f_2 = a_1 x^2 + a_2 y^2 + a_3 z^2.$$

However, care must be taken because as functions on the unit sphere, x^2, y^2, and z^2 are not orthonormal. By fixing an orthonormal basis $\{v_1, v_2, v_2\}$ for the span of $\{x^2, y^2, z^2\}$

we can replace the harmonic representation $\text{SH}(f)$ defined in Section 3.2 with the more discriminating representation:

$$\text{SHQ}(f) = \{R^{-1}(a_1, a_2, a_3), \|f_1\|, \|f_3\|, \ldots\}$$

where R is the matrix whose columns are the orthonormal vectors v_i.

4. Properties of the Spherical Harmonic Representation

This section provides a mathematical analysis of some of the properties and limitations of the spherical harmonic representation. In particular, we describe how the similarity of spherical descriptors, defined as the optimum over all rotations, relates to the similarity of their harmonic representations. We also describe the way in which information is lost in going from a spherical shape descriptor to its harmonic representation.

1. **Similarity**: The L_2-difference between the harmonic representations of two spherical functions is a lower bound for the minimum of the L_2-difference between the two functions, taken over all possible orientations. To see this, we let $f(\theta, \phi)$ and $g(\theta, \phi)$ be two spherical functions, and observe that:

$$\begin{aligned}
\|\text{SH}(f) - \text{SH}(g)\|^2 &= \sum_{l=0}^{\infty} (\|f_l\| - \|g_l\|)^2 \\
&\leq \sum_{l=0}^{\infty} (\|f_l - g_l\|)^2 = \|f(\theta, \phi) - g(\theta, \phi)\|^2.
\end{aligned}$$

Similarly, if we consider the rotation invariant representation described in Section 3.3, we get:

$$\begin{aligned}
\|\text{SH}(f) - \text{SH}(g)\| &\leq \|\text{SHQ}(f) - \text{SHQ}(g)\| \\
&\leq \|f - g\|.
\end{aligned}$$

But as we have shown, the harmonic representations are invariant to rotation, so we get:

$$\begin{aligned}
\|\text{SH}(f) - \text{SH}(g)\| &\leq \|\text{SHQ}(f) - \text{SHQ}(g)\| \\
&\leq \min_{R \in SO(3)} \|f - R(g)\|.
\end{aligned}$$

2. **Information Loss**: In general, if a spherical function $f(\theta, \phi)$ is band-limited with bandwidth b, then we can express f as:

$$f(\theta, \phi) = \sum_{l=0}^{b} \sum_{m=-l}^{l} a_{lm} Y_l^m(\theta, \phi).$$

Thus, the space of spherical functions with bandwidth b is of dimension $O(b^2)$. The harmonic representation, however, is of dimension $O(b)$ so that a full dimension worth of information is lost in going from a spherical function to its harmonic representation. This information loss happens in two different ways:

- First, we treat the different frequency components independently. Thus if we write:

$$f = \sum_{l=0}^{b} f_l \quad \text{and} \quad g = \sum_{l=0}^{b} R_l(f_l)$$

where R_l are rotations, then the descriptors of the functions f and g will be the same. That is, the descriptor is unchanged if we apply different rotations to the different frequency components of a spherical function. Figure 2 shows a visualization of this for two spherical functions. The one on the bottom is obtained from the one on the top by applying a rotation to only one of the frequency components. Though the two functions differ by more than a single rotation, there spherical harmonic descriptors are the same. (An analagous form of information loss occurs with Fourier Descriptors where the phases of different frequencies are discarded independently.)
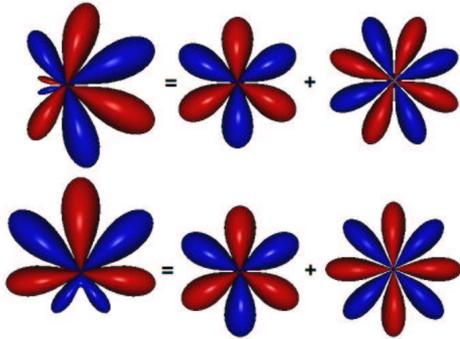
Figure 2: *The bottom spherical function is obtained by rotating one of the frequency components of the top one. Despite the fact that there is no rotation transforming the function on the top to the one on the bottom, the descriptors of the two functions are the same.*

- Second, for each frequency component f_l, the harmonic representation only stores the energy in that component. For $l > 2$ it is *not* true that if $\|f\| = \|g\|$ then there is a rotation R such that $R(f) = g$. Thus knowing only the norm of the l-th frequency component does not provide enough information to reconstruct the component up to rotation. (This form of information loss does not occur with Fourier Descriptors, as two circular functions with the same amplitude and frequency can only differ by phase/rotation.) Figure 3 shows a visualization of this for three spherical functions. The functions are all of the same frequency and have the same amplitude, but there is no rotation that can be applied to transform them into each other.
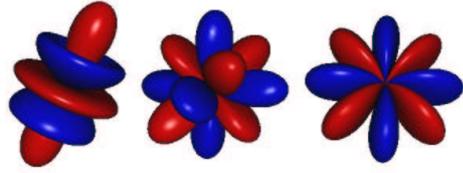
Figure 3: *Three spherical functions of the same amplitude and frequency are shown. Note that there is no rotation transforming any one of them into the other.*

5. Extensions to Voxel Descriptors

In Section 3 we presented a method for obtaining rotation invariant representations of spherical functions. In this section we show how this method can be generalized to obtain rotation invariant representations of voxel descriptors.

5.1. Rotation Invariant Representations

In order to obtain a rotation invariant representation of a voxel grid we use the obsevation that rotations fix the distance of a point from the origin. Thus, we can restrict the voxel grid to concentric spheres of different radii, and obtain the spherical harmonic representation of each spherical restriction independently. This process is demonstrated in Figure 4: First, we restrict the voxel grid to a collection of concentric spheres. Then, we represent each spherical restriction in terms of its frequency decomposition. Finally, we compute the norm of each frequency component, at each radius. The resultant rotation invariant representation is a 2D grid indexed by radius and frequency.

5.2. Properties

In addition to the information loss described in Section 4, the method described above also loses information as a result of the fact that the representation is invariant to independent rotations of the different spherical functions. For example, the plane in Figure 5 (right) is obtained from the one on the left by applying a rotation to the interior part of the model. While the two models are not rotations of each other, the descriptors obtained are the same.

6. Experimental Results

To measure the efficacy of the spherical harmonic representation, we computed a number of spherical shape descriptors, and compared matching results when the spherical functions were aligned by PCA with the results obtained when the spherical harmonic representation was used.

In order to evaluate our method we computed the following spherical descriptors:

- **Extended Gaussian Image** [4]: This is a description of a surface obtained by binning surface normals.
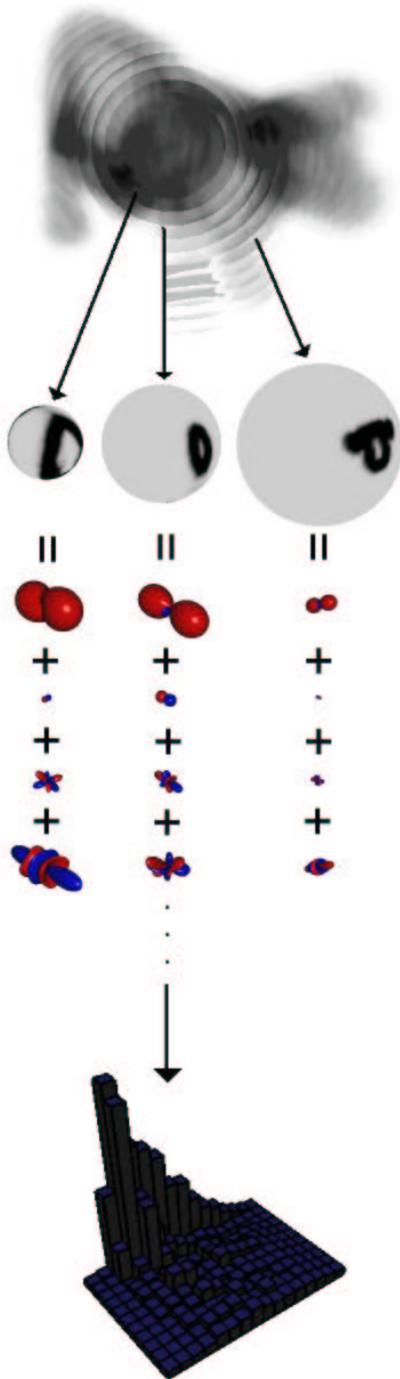
Figure 4: *We compute a rotation invariant descriptor of a voxel grid by intersecting the model with concentric spheres, computing the frequency decomposition of each spherical function, and computing the norms of each frequency component at each radius. The resultant rotation invariant representation is a 2D grid indexed by radius and frequency.*

Figure 5: *The model on the right is obtained by applying a rotation to the interior part of the model on the left. While the models differ by more than a single rotation, their rotation invariant representations are the same.*

- **Radial Distribution**: This is a description of a surface that associates to every ray through the origin, the average distance and standard deviation of points on the intersection of the surface with the ray.
- **Spherical Extent Function** [6]: This is a description of a surface associating to each ray from the origin, the value equal to the distance to the last point of intersection of the model with the ray.
- **Sectors**: This is a description of a surface associating to each ray from the origin, the amount of surface area that sits over it. This is a continuous implimentation of the shells in Shape Histograms [5], with sectors chosen to correspond to a single cell within the 64×64 representation of the sphere.
- **Shape Histogram** [5]: This a finer resolution of the Sector descriptor that breaks up the bounding sphere of the model into a collection of shells and computes the sector descriptor for the intersection of the model with each one.
- **Voxel** [12]: This is a description of a shape as a voxel grid, where the value at each point is given by the negatively exponentiated Euclidean Distance Transform of the surface.

We evaluated the performance of each method by testing how well they classified models within a test database. The database consisted of 1890 "household" objects provided by Viewpoint [17]. The objects were clustered into 85 classes, based on functional similarities, largely following the groupings provided by Viewpoint and classes ranged in size from 5 models to 153 models, with 610 models that did not fit into any meaningful classes [12]. Classification performance was measured using precision/recall plots, which which gives the percentage of retrieved information that is relevant as a function of the percentage of relevant information retrieved.

We computed the spherical representations as 64×64 grids corresponding to regular sampling along the lines of longitude and lattitude and we used SPharmonicKit 2.5 [18] to obtain the spherical harmonic representation as an array of 33 floating point numbers. Both the spherical descriptors and their spherical harmonic representations were compared

using the L_2-difference. The results of the classification experiment are show in in Figure 6.
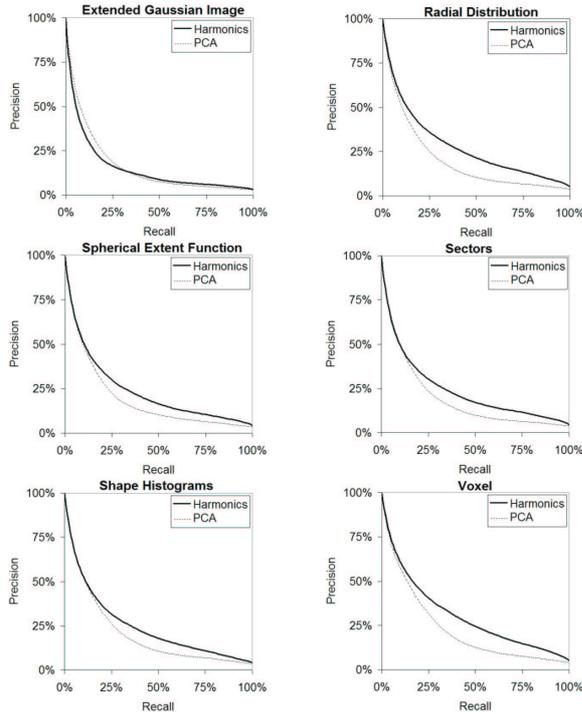
Figure 6: *Precision vs. Recall plots comparing the performance of aligned spherical descriptors with the performance of their harmonic representations. Note that for most of the representations the harmonic descriptor outperforms the canonically aligned one.*

As the results indicate, the application of the Spherical Harmonic Representation improves the performance of most of the descriptors. The improvement of the matching results is particularly meaningful when we consider the fact that the Spherical Harmonic Representation reduces a 2D descriptor into a 1D array of energy values. Thus, the representation not only provides better performance, but it does so with fewer bits of information.

7. Discussion

In this section we present a discussion of the results in Section 6. In particular, we analyze the case of the Extended Gaussian Image, and discuss how this reflects on the general limitations of the Spherical Harmonic Representation. We also evaluate the implications of the Spherical Harmonic Representation for database retrieval.

7.1. Limitations

The analysis described in Appendix A provides a mathematical interpretation of the failing of PCA-alignment. This analysis makes the assumption that we are looking at the

general class of spherical functions, so that frequency components align independently. However, in certain shape applications this may not be the case and the descriptors obtained may fall into a restrictive subset of spherical functions. In these cases it is possible that the alignment of different frequency components are correlated and PCA-alignment performs well.

Such a case may occur when the spherical functions are primarily axis aligned, so that, up to rotation, they can be described as:

$$\sum a_k x^k + b_k y^k + c_k z^k$$

and the alignments of the different frequency components are strongly correlated. This is the case for the Extended Gaussian Image [4] which describes a polygonal model by the distribution of normal vectors over the unit sphere. When the database of models is restricted to household objects, the obtained descriptors are primarily axis aligned (see Figure 7) and principal axis alignment may provide optimal alignment, (as indicated by the improved performance in Figure 6).

Figure 7: *Images of models of a vase, a chair, and scissors, with their associated Extended Gaussian Images. Note that the EGIs are mainly axial functions and consequently are well aligned by PCA.*

7.2. Implications for Model Databases

Much of the research presented in this paper is guided by the increased proliferation and accessibility of 3D models. These models have been gathered into databases, and one of the challenges has been to design matching implementations that are well suited for database retrieval. The spherical harmonic representation presented in this paper addresses this challenge in two ways:

1. While a spherical function of bandwidth b requires $O(b^2)$ space, its spherical harmonic representation is of size $O(b)$. Consequently, the spherical harmonic representations provide a more compact representation of the descriptors, and can be compared more efficiently. (For each method compared in Section 6, Table 2 shows the space requirements of the descriptor and its Spherical Harmonic Representation.)
2. Furthermore, the Spherical Harmonic Representations are based on a frequency decompostion of a spherical function. Consequently, the representation is inherently multiresolutional and this property can be used to guide indexing schemes for efficient retrieval.

Representation	PCA-Aligned	Harmonic
EGI	64×64	33
Spherical Extent Function	64×64	33
Radial Distribution	$2 \times 64 \times 64$	2×33
Sectors	64×64	33
Shape Histogram	$4 \times 64 \times 64$	4×33
Voxel	$32 \times 64 \times 64$	32×33

Table 2: *The number of floating point numbers used to describe each representation. This table demonstrates that the Spherical Harmonic Representation provides a representation that reduces the dimensionality of the space required for storing the descriptor.*

8. Conclusion and Future Work

In this paper we have introduced the Spherical Harmonic Representation, a rotation invariant representation of spherical functions in terms of the energies at different frequencies. We have shown that this representation provides a method for improving the performance of many canonically aligned spherical descriptors in tasks of shape matching. In addition to providing better matching performance this rotation invariant representation also reduces the dimensionality of the existing descriptors improving both the time and space requirements of these methods.

This work suggests a number of challenges that we would like to consider in the future: First, we would like to explore the possibility of generalizing this method to voxel grids using Zernike moments. Second, we would like to consider methods for reducing the rotation independence of the different frequency components, and, in the case of voxel grids, of the different radial components. Finally, we would like to explore extending this method to capture more rotation invariant information in the higher frequency components, allowing us to truly reconstruct each frequency component uniquely up to rotation.

References

1. Delingette, H., Hebert, M., Ikeuchi, K.: A spherical representation for the recognition of curved objects. In: Proc. ICCV. (1993) 103–112

2. Besl, P.: Triangles as a primary representation. Object Recognition in Computer Vision **LNCS 994** (1995) 191–206

3. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Matching 3d models with shape distributions. Shape Matching International (2001) 154–166

4. B. Horn, B.: Extended gaussian images. PIEEE **72** (1984) 1656–1678

5. Ankerst, M., Kastenmüller, G., Kriegel, H.P., Seidl, T.: 3D shape histograms for similarity search and classification in spatial databases. In: Proc. SSD. (1999)

6. Vranic, D., Saupe, D.: 3d model retrieval with spherical harmonics and moments. Proceedings of the DAGM (2001) 392–397

7. Gain, J., Scott, J.: Fast polygon mesh querying by example. SIGGRAPH Technical Sketches (1999) 241

8. Kazhdan, M., Chazelle, B., Dobkin, D., Finkelstein, A., Funkhouser, T.: A reflective symmetry descriptor. European Conference on Computer Vision (ECCV) (2002) 642–656

9. Elad, M., Tal, A., Ar, S.: Content based retrieval of vrml objects - an iterative and interactive approach. EG Multimedia (2001) 97–108

10. Horn, B., Hilden, H., Negahdaripour, S.: Closed form solution of absolute orientation using orthonormal matrices. Journal of the Optical Society (1988) 1127–1135

11. Zahn, C., Roskies, R.: Fourier descriptors for plane closed curves. IEE Transaction on Computers **21** (1972) 269–281

12. Funkhouser, T., Min, P., Kazhdan, M., Chen, J., Halderman, A., Dobkin, D., Jacobs, D.: A search engine for 3d models. ACM TOG (2003) 83–105

13. Besl, P., McKay, N.: A method for registration of 3d shapes. IEEE TPAMI (1992) 239–256

14. Zhang, Z.: Iterative point matching for registration of free-form curves and surfaces. IJCV (1994) 119–152

15. Ballard, D.: Generalized hough transform to detect arbitrary patterns. IEEE PAMI **12** (1981) 111–122

16. Lamdan, Y., Wolfson, H.: Geometric hashing: A general and efficient model-based recognition scheme. Proceedings of the 2nd International Conference on Computer Vision (1988) 238–249

17. Labs, V.D.: http://www.viewpoint.com (2001)

18. 2.5, S.: http://www.cs.dartmouth.edu/ gee-long/sphere/ (1998)

Appendix A: A Signal Processing Framework for PCA

This appendix presents a signal processing framework for analyzing the implications and limitations of model alignment via PCA. We define a spherical function characterizing the radial variance of a shape along different rays from the origin. In particular, for a model S and a direction v we set:

$$\text{RV}(S, v) = \lim_{\alpha \to 0} \int_{C(v, \alpha) \cap S} \frac{\|x\|^2}{2\pi(1 - \cos(\alpha))} \, dx$$

where $C(v, \alpha)$ is the cone with apex at the origin, angle α and direction v, and $2\pi(1 - \cos(\alpha))$ is the area of the intersection of the cone with the unit sphere (see Figure 8). That is, $\text{RV}(S, v)$ gives the sum of the square of the distances of the points lying on the intersection with S and the ray, from the origin, with direction v. Figure 9 shows a visualization of the *Radial Variance* for a cube by scaling the radius of each point on a sphere in proportion to the value of the function at that point. Note that the function scales the points at the corners of the cube more drastically because: (1) we integrate
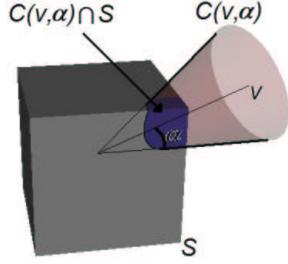
Figure 8: *The value of the* Radial Variance *in the direction v is defined by intersecting the model with a cone, in the direction v, with small angle* α, *and integrating the square of the distance over the intersection of the model with the cone.*

Figure 9: *The* Radial Variance *can be visualized by displacing the radius at a point on the sphere, in proportion to the value of the function at that point.*

the *square* of the distance to the origin over each patch, and (2) the angle between the point on the sphere and the surface normal is large, so that more surface area projects onto a spherical patch.

What is valuable about this function is that for any surface S, the function has the property that:

$$\int_S x_i x_j \, dx = \langle RV(S, v), x_i x_j \rangle_{S^2}.$$

That is, the second (and 0-th) order components of the radial variance are precisely the terms of the covariance matrix of the model. This function gives a representation of the initial model in a signal processing framework that allows us to make two observations:

1. Because of the orthogonality of the frequency components, principal axis registration does not take into account information at non second-order frequencies and hence makes no guarantees as to how they align.
2. Aligning two models using their principal axes provides the optimal alignment for their second order components, as will be shown in the following theorem:

Theorem: If f and g are two spherical functions consisting of only constant and second order harmonics, then the L_2-difference between the two is minimized when each is aligned to its own principal axes.

Proof: Because f and g consist of only constant and second order terms, we can represent the functions by symmetric matrices A and B where

$$f(v) = v^t A v \qquad \text{and} \qquad g(v) = v^t B v.$$

If we assume that A and B are already aligned to their principal axes we get:

$$A = \begin{pmatrix} a_1 & 0 & 0 \\ 0 & a_2 & 0 \\ 0 & 0 & a_3 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_1 & 0 & 0 \\ 0 & b_2 & 0 \\ 0 & 0 & b_3 \end{pmatrix}$$

Thus, if R is any rotation we get:

$$\langle R^t(f), g \rangle = (\alpha - \beta) \text{Trace}(ARBR^t) + \beta \sum_{i,j=1}^{3} a_i b_j$$

where $\alpha = \int_{S^2} x^4 dx$ and $\beta = \int_{S^2} x^2 y^2 dx$ define the lengths and angles between the functions x_i^2 on the unit sphere. We would like to show that the dot product is maximized when R is a permutation matrix so that RAR^t is diagonal.

Using the fact that the differentials of a rotation R are defined by RS where S is a skew-symmetric matrix, it suffices to solve for:

$$0 = \frac{d}{dt}\Big|_{t=0} \text{Trace}(A(R+tRS)B(R^t - tSR^t))$$
$$= \text{Trace}(R^t AR(SB - BS))$$

But S is a skew-symmetric matrix so that, $SB - BS$ is a symmetric matrix with 0's along the diagonal:

$$SB - BS = \begin{pmatrix} 0 & (b_2 - b_1)S_{12} & (b_3 - b_1)S_{13} \\ (b_2 - b_1)S_{12} & 0 & (b_3 - b_2)S_{23} \\ (b_3 - b_1)S_{13} & (b_3 - b_2)S_{23} & 0 \end{pmatrix}$$

Thus, if $R^t AR$ is a diagonal matrix then the derivative is zero, independent of the choice of S. Conversely, if the b_i are distinct and $R^t AR$ is not diagonal, we can always choose values for S_{12}, S_{13}, and S_{23} such that the derivative is non-zero, implying that if $R^t AR$ is not diagonal it cannot maximize the dot product. (Note that if $b1 = b2 = b3$ then B is a constant multiple of the identity so that the dot product is independent of the choice of rotation. Similarly, if $b_i = b_j$ then rotations in the plane spanned by x_i and x_j also do not change the dot product.)

This shows that the L_2-difference between f and g is at an extremum if and only if A and B are both diagonal matrices. The minimum L_2-difference is then attained when $\sum a_i b_i$ is maximal. So that if $a_1 \geq a_2 \geq a_3$ then we must also have $b_1 \geq b_2 \geq b_3$, and the L_2-difference between f and g is minimized precisely when f and g are aligned to their principal axes.