Harmonic 3D Shape Matching

Michael Kazhdan, and Thomas Funkhouser *
Princeton University

1 Introduction

With the advent of the world wide web, the number of available 3D models has increased substantially and the challenge has changed from "How do we generate 3D models?" to "How do we find them?" In this sketch we describe a new 3D model matching and indexing algorithm that uses spherical harmonics to compute discriminating similarity measures without requiring repair of model degeneracies or alignment of orientations. It provides 46–245% better performance than related shape matching methods during precision-recall experiments, and it is fast enough to return query results from a repository of 20,000 models in under half a second.

2 Model Representation and Matching

The main challenge in designing a model matching algorithm is to find a computational representation of shape for which an index can be built and geometric matching can be performed efficiently. Generally speaking, the following properties are desirable for a shape representation. It should be: (1) quick to compute, (2) concise to store, (3) easy to index, (4) invariant under similarity transforms, and (5) independent of 3D object representation, tessellation, genus, or topology.

Previous work in defining shape representations for matching can be classified into two broad categories. The first class of representations use methods such as PCA to align the model into a canonical coordinate frame and then define the shape representation with respect to this orientation. Such methods include, among others, Moments [Elad et al. 2001] and Extended Gaussian Images [Horn 1984]. The second class of methods define representations that are invariant under rotation and include methods such as Shape Histograms [Ankerst et al. 1999] and Shape Distributions [Osada et al. 2001]. In our experiments we have found that PCA based methods are unstable as a result of multiplicity of eigenvalues and sensitivity to outliers. The shape representation that we have designed is both rotation invariant and discriminating, characterizing a shape in terms of the clustering of mass on different concentric spheres.

The steps for computing the harmonic shape representation are outlined in Figure 1: (1) Given a model, we rasterize its polygons into a $64 \times 64 \times 64$ voxel grid, (assigning a voxel a value of 1 if it was within one voxel of a point on the boundary, and a value of 0 otherwise). The model is aligned so that its center of mass is at the center of the grid, and so that its bounding sphere has radius 32. (2) Treating it as a function defined in three-space, we decompose the voxel grid into 32 spherical functions by restricting the voxel grid to spheres with radii 1 through 32. (3) We decompose each of these functions as a sum of its first 16 harmonic components, analogous to a Fourier decomposition into different frequencies. (4) Using the fact that rotations do not change the norm of the harmonic components, we define the signature of each spherical function as a list of these 16 norms. (5) Finally, we combine these different signatures to obtain a 32×16 signature for the 3D model. The resultant rotation invariant signature is a two-dimensional grid where the value of the (i, j)-th index is equal to the norm of the j-th order component of the spherical function on the sphere of radius i.

To compare two harmonic representations, we simply compute the Euclidean distance between them. Thus, finding the K closest models to a query is equivalent to solving the nearest-neighbor problem in a 32×16 dimensional space. Although this problem is known to be hard in the worst case, we can build a search algorithm that works efficiently in practice by taking advantage of the multiresolution nature of the harmonic decomposition to guide a dimension reduction indexing method.

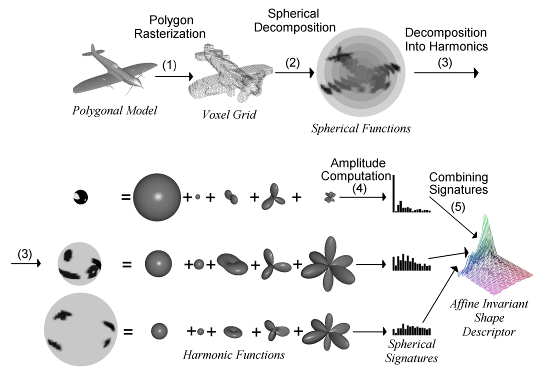

Figure 1: Computing the Harmonic Shape Representation

3 Results

We compared the shape classification performance of our method to five existing methods, using a test database of 1890 models provided by Viewpoint. The models were clustered into 85 classes based on functional similarities. The smallest class had 5 models, the largest 143, and 610 models did not fit into any meaningful class. The other methods tested were: Moments, Extended Gaussian Images, Shape Histograms, and D2 Shape Distributions. The figure on the right shows the precision vs. recall graphs for each method. Note that our method has precision values on average 42 % higher than D2, 60 % higher than Shape Histograms, 126 % higher than EGIs, and 245 % higher than moments.

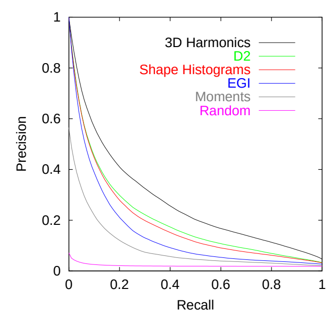

4 Conclusion

In summary, this sketch investigates a new method for model matching and indexing. The main research contribution is a new shape representation that allows for efficient, robust, and discriminating querrying of models in large databases.

References

Ankerst, M., Kastenmüller, G., Kriegel, H.-P., and Seidl, T. 1999. 3d shape histograms for similarity search and classification in spatial databases. In *Proc. SSD*, 207–226.

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

Horn, B. 1984. Extended gaussian images. *Proc. of the IEEE 72*, 12 (December), 1671–1686.

Osada, R., Funkhouser, T., Chazelle, B., and Dobkin, D. 2001. Matching 3d models with shape distributions. *Shape Modeling International* (May).

*e-mail: mkazhdan, funk @cs.princeton.edu