Symmetry Descriptors and 3D Shape Matching

Michael Kazhdan, Thomas Funkhouser and Szymon Rusinkiewicz

Department of Computer Science, Princeton University, Princeton NJ

Abstract

In this paper, we present the Symmetry Descriptors of a 3D model. This is a collection of spherical functions that describes the measure of a model's rotational and reflective symmetry with respect to every axis passing through the center of mass. We show that Symmetry Descriptors can be computed efficiently using fast signal processing techniques, and demonstrate the empirical value of Symmetry Descriptors by showing that they improve matching performance in a variety of shape retrieval experiments.

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

1. Introduction

Symmetry has long been recognized as playing an integral role in human recognition [Att55, Vet92]. It is characteristic of repeating patterns within a model, and can be used to guide reconstruction, compression, and classification. This awareness has motivated the development of a wide range of techniques for identifying the symmetries of a 2D image [Ata85, Wol85, Hig86, Sun95, Mar89]. However, the increased complexity of the rotation group in three-dimensions has resulted in little research on symmetry detection in 3D. Methods for measuring individual symmetries have been proposed [Zab94, Zab95], and a general approach for characterizing the measure of all reflective symmetries has been described [Kaz02, Kaz04], but no analogous method for describing all rotational symmetries exists.

In this paper, we present the *Symmetry Descriptors* of a model. This is a generalization of the Reflective Symmetry Descriptor presented in [Kaz02, Kaz04]. It represents a 3D model as a collection of spherical functions that give the measure of a model's reflective and rotational symmetry, with respect to every axis passing through the center of mass. Thus, it can be used not only to identify axes of perfect symmetry, but also to measure the quality of symmetry with respect to any axis. Specifically, the measure of k-fold symmetry of a model around some axis is defined to be the magnitude of the projection of the model onto the space of models having that symmetry.

Figure 1 shows a visualization of the Symmetry Descriptors of two models. The descriptors are represented by scaling points on the unit sphere in proportion to the measure of symmetry, so that points corresponding to axes of near symmetry are pushed out from the origin and points corresponding to axes of near anti-symmetry are pulled in to the origin. Thus, for the 2-fold (respectively k-fold) symmetry descriptors, peaks in the descriptors correspond to axes of near perfect 2-fold (respectively k-fold) rotational symmetry. Similarly, for the reflective symmetry descriptors, peaks correspond to unit vectors perpendicular to planes of near perfect reflective symmetry.

The contribution of our work is three-fold. First, we define a continuous measure for the reflective and rotational symmetry of a 3D model. Second, we provide an efficient algorithm for computing the measure for all symmetries about a model's center of mass. Third, we present experimental results evaluating the empirical value of the symmetry descriptors in shape retrieval applications. In these experiments, we find that symmetry can be used to augment existing methods for matching 3D shapes, providing enhanced discrimination and matching performance without sacrificing efficiency.
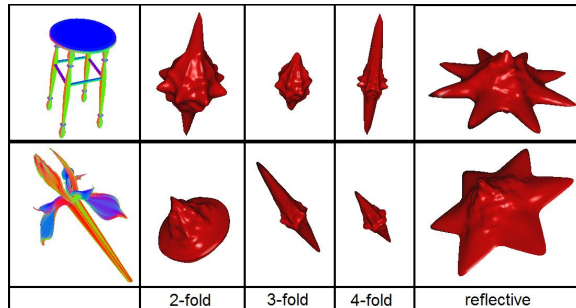
| | 2-fold | 3-fold | 4-fold | reflective |

Figure 1: *A visualization of the symmetry descriptors for a stool and an iris. The visualization is obtained by scaling unit vectors on the sphere in proportion to the measure of rotational symmetry about the axis through the center of mass, in the direction of the vector, and the measure of reflective symmetry about the plane through the center of mass, normal to the vector.*

The rest of this paper is structured as follows. Section 2 reviews related work in the area of symmetry detection. Section 3 provides a theoretical overview of symmetry, and defines the *Symmetry Descriptors*. Section 4 describes an efficient method for computing the Symmetry Descriptors of spherical and voxel representations of a 3D model, while Section 5 summarizes some properties of the descriptors. In Section 6, we describe how symmetry information can be incorporated into existing rotation invariant representations, and in Section 7, we evaluate the contribution of symmetry augmentation in experiments comparing the retrieval performance of the original representation with the retrieval performance of the augmented representation. Finally, we conclude in Section 8 by summarizing our work.

2. Related Work

Early approaches to symmetry detection focused on detecting the symmetries of planar point sets [Ata85, Wol85, Hig86]. These methods reduced the symmetry detection problem to a detection of symmetry in circular strings, and used efficient substring algorithms (e.g., [Knu77]) to detect the symmetries by searching for the appearance of a string within its concatenation with itself. While these methods had the theoretical advantage of efficiently evaluating all possible symmetries, they were impractical in empirical settings since they were algorithms that could only identify the perfect symmetries of a model. Thus if a symmetric model had even a small amount of noise, these methods would fail to identify its symmetries.

In order to address this issue, Zabrodsky et al. [Zab94, Zab95] defined a continuous measure of symmetry which transformed the binary question: "Does a model have a given symmetry?" to the continuous question: "How much of a given symmetry does a model have?" The measure of symmetry was defined as the minimum amount

of work needed to transform a model into a symmetric model, measured as the sum of the squares of the distances that points would need to be moved. This approach made it possible to evaluate symmetries in the presence of noise, but suffered from the fact that it depended on the establishment of point correspondences. While this issue could be addressed in the case of 2D curves with uniform sampling, it made it difficult to generalize the method to 3D where uniformly sampling surfaces is often impossible.

The difficulty of establishing point correspondences for matching surfaces in 3D has motivated the development of *shape descriptors* which represent a 3D model by a function defined on a canonical domain, independent of the initial model's shape or topology. (For a general review of such methods see [Pop94, Tan04].) For these descriptors, matching two models could now be performed without explicitly establishing correspondences, by comparing the values of the corresponding shape descriptors at each point.

The advantage of the canonical parameterization of shape descriptors was leveraged in a number of symmetry detection algorithms [Oma96, Sun97]. These methods used the fact that the covariance ellipsoid of a 3D model rotates with the model, so that a model could only have symmetries where its covariance ellipsoid had them. Since the only axes of symmetry of an ellipsoid have to align with its principal axes, this provided an efficient way to identify candidate axes of symmetry. The actual quality of an axis as an axis of symmetry would then be measured by comparing the shape descriptor of the model with the shape descriptors of the rotations and reflections of the model about the candidate axis. This method had the advantage of providing a continuous measure of symmetry for candidate axes of symmetry without necessitating the establishment of point correspondences. Furthermore, the method was a general one that could be applied to wide class of shape descriptors. However, the method's dependence on PCA for the identification of candidate axes could only guarantee the correct identification of symmetry axes for models with perfect symmetry.

Motivated by the ease of evaluating symmetry using shape descriptors, and the efficiency of exhaustive search provided by early substring matching approaches, efficient methods for evaluating the symmetries of a 2D model, at every symmetry, were developed. The key idea of these approaches was the generalization of discrete substring matching to continuous correlation with the Fast Fourier Transform. These methods [Sun95, Mar89] compute the symmetries of a model by using correlation to compare the shape descriptor of a 2D model with all of its rotations and reflections. This approach was a general one that could be applied to any shape descriptor that represented a model with a function defined either on a circle, or in 2D.

The dependence of these methods on the FFT made them hard to generalize to shape descriptors that represented a 3D model with either a spherical function or a function in

3D. In [Kaz02, Kaz04] a method is described for computing the measure of reflective symmetries for all planes passing through the origin. For a spherical descriptor of size $O(N^2)$ (respectively 3D function of size $O(N^3)$) the method computes the measures of reflective symmetry in $O(N^3 \log N)$ (respectively $O(N^4 \log N)$) time. The efficiency of this approach relies on the use of the FFT to compute correlation with respect to a single axis efficiently and a generalization of this approach to general symmetry detection would result in algorithms that have complexity $O(N^4 \log N)$ for spherical functions and $O(N^5 \log N)$ for 3D functions.

In this work, we show how the analogs of the FFT and inverse FFT on the sphere, namely the Fast Harmonic Transform and the Fast Inverse Wigner-D Transform, can be used to compute the measure of all symmetries efficiently. In particular, we describe a method for computing the measure of all reflective and rotational symmetries of both spherical functions and 3D functions in $O(N^4)$ time, providing a method for computing all symmetries of a model about its center of mass, in less time than previous methods required to compute only the reflective symmetries.

3. Measuring Symmetry

The first issue we must address is to define a measure of symmetry for a 3D model with respect to an axis of k-fold rotation or a plane of reflection. To this end, we describe a method for computing the symmetries of any shape descriptor that represents a 3D model as a function defined either on the sphere or in 3D. We begin by describing a general approach for measuring symmetry, and then present the implications for measuring the symmetries of a 3D model.

3.1. Symmetry Detection

Definition: Given a vector space V and a group G that acts on V, we say that $v \in V$ is *symmetric with respect to G* if $\gamma(v) = v$ for all $\gamma \in G$.

Definition: We define the *symmetry distance* of a vector v with respect to a group G as the L_2-distance to the nearest vector that is symmetric with respect to G:

$$\mathrm{sd}_G(v) = \min_{w|G(w)=w} \|v - w\|.$$

Using the fact that the vectors that are invariant to G define a subspace of V, it follows that the nearest G-invariant vector w is the projection of v onto the subspace of invariant vectors. That is, if we define π_G to be the projection onto the subspace invariant under the action of G and we define $\pi_G^\perp$ to be the projection onto the orthogonal subspace then:

$$\mathrm{sd}_G(v) = \|v - \pi_G(v)\| = \|\pi_G^\perp(v)\|$$

so that the symmetry distance of v with respect to G is the length of the projection of v onto a subspace indexed by G.

In general, computing the projection of v onto the subspace of vectors invariant under the action of G is a difficult task. However, in our case we can use the fact that the elements of G are orthogonal transformations. In particular, we can apply a theorem from representation theory [Ser77] stating that a projection of a vector onto the subspace invariant under the action of an orthogonal group is the average of the vector over the different elements in the group. Thus, in the case of a vector v and a group G, we get:

$$\mathrm{sd}_G^2(v) = \left\| v - \frac{1}{|G|} \sum_{\gamma \in G} \gamma(v) \right\|^2 = \|v\|^2 - \frac{1}{|G|} \sum_{\gamma \in G} \langle v, \gamma(v) \rangle$$

giving an expression for the symmetry distance in terms of the length of v and the dot-products of v with its image under the action of G.

3.2. Symmetry Descriptors in 3D

In order to evaluate the measure of symmetry of a 3D model, it is necessary to compare a model with its reflections/rotations. A variety of shape descriptors can be used to compare the model with its transformation, and in this paper we focus on those that represent a model by a spherical, or 3D, function that rotates with the model.

Notation: For any integer k and any unit vector p we let G_p^k denote the *k-fold rotational symmetry group with respect to p*. If k is positive, then G_p^k is the group generated by the transformation $r_p^{2\pi/k}$ which is the rotation about the axis p by the angle $2\pi/k$. If k is negative, then G_p^k is the group generated by the transformation $r_p^{2\pi/k} \cdot A$, where A is the antipodal map, sending a point q to the point $-q$.

For example, $G_{(1,0,0)}^3$ is the group generated by rotating by $120°$ about the x-axis, consisting of 3 elements, while $G_{(0,1,0)}^{-2}$ is the group generated by reflecting through the xz-plane, consisting of two elements.

Definition: Given a shape descriptor f, we define its *k-fold symmetry descriptor* as the function on the sphere whose value at a point p describes the amount of f that is symmetric with respect to G_p^k and the amount of f that is anti-symmetric:

$$\mathrm{SD}_k(f, p) = \left(\|\pi_{G_p^k}(f)\|, \|\pi_{G_p^k}^\perp(f)\| \right)$$

where $\pi_{G_p^k}$ is the projection onto the space of functions that are k-fold symmetric about the axis p, and $\pi_{G_p^k}^\perp$ is the projection onto the orthogonal complement. (Note that since $\|f\|^2 = \|\pi_{G_p^k}(f)\|^2 + \|\pi_{G_p^k}^\perp(f)\|^2$ it suffices to compute one of $\|\pi_{G_p^k}(f)\|$ and $\|\pi_{G_p^k}^\perp(f)\|$. Despite the redundancy, we store both values, as they can be used for bounding shape similarity.)

4. Computing the Symmetry Descriptors

We will now show how to compute all the *k*-fold symmetry descriptors of a shape descriptor efficiently. We begin by describing the method for shape descriptors that represent a model with a spherical function. Then, we generalize the method to shape descriptors that represent a model with a function defined in 3D. For both cases, the key idea is that in computing the symmetry descriptors of a shape descriptor f, it is necessary to compare f with its rotations. This amounts to computing the autocorrelation of f across the group of rotations, $\langle f, \gamma(f) \rangle$, and, while a brute force algorithm would be prohibitively slow, we show how spherical signal processing can be used to compute the autocorrelation efficiently.

4.1. Efficient Spherical Autocorrelation

We begin by describing an efficient method for computing the autocorrelation of a spherical function f across the space of rotations. We express f in terms of its spherical harmonic decomposition:

$$f = \sum_{l=0}^{b} \sum_{|m| \leq l} a_{l,m} Y_l^m$$

where b is the band width of the function f. This decomposition has the property that for any rotation γ we have:

$$\langle Y_l^m, \gamma(Y_{l'}^{m'}) \rangle = 0 \qquad \forall l \neq l'.$$

This allows us to compute the autocorrelation of f by cross multiplying spherical harmonic coefficients within each frequency l, ignoring cross-frequency terms:

$$\langle f, \gamma(f) \rangle = \sum_{l=0}^{b} \sum_{|m|,|m'| \leq l} a_{l,m} \overline{a_{l,m'}} \langle Y_l^m, \gamma(Y_l^{m'}) \rangle.$$

This gives an expression of the function $\langle f, \gamma(f) \rangle$ as a linear sum of the functions $\langle Y_l^m, \gamma(Y_l^{m'}) \rangle$. Since these are precisely the Wigner-D functions, a fast inverse Wigner-D transform [Kos03, Sof03] gives the autocorrelation of f over the space of rotations.

Since the spherical harmonic decomposition can be computed in $< O(b^3)$ time, the cross multiplication of harmonic coefficients takes $O(b^3)$ time, and since the inverse Wigner-D transform can be done in $\leq O(b^4)$ time, we compute the autocorrelation of f across all rotations in time $\leq O(b^4)$.

4.2. Efficient 3D Autocorrelation

In order to compute the symmetry descriptors of a 3D function, we decompose the 3D function as a collection of spherical functions at different radii and use the method described above to compute the symmetry descriptors of the collection of spherical functions. In particular, given a function f defined for all points $|x| \leq 1$, we set f_r to be the restriction of the function f to the sphere with radius r:

$$f_r(p) = f(rp)\sqrt{4\pi r^2}$$

where p is a unit vector and $\sqrt{4\pi r^2}$ is the change of variable term accounting for the area of the sphere with radius r.

Discretely sampling the different radii, we can express f as a collection of spherical functions $\{f_1, \ldots, f_N\}$, with $N = O(b)$ where b is the sampling rate. Expressing each of these in terms of its spherical harmonic representation, we get:

$$f_i = \sum_{l=0}^{b} \sum_{|m| \leq l} a_{l,m}[i] Y_l^m$$

and for any rotation γ we have

$$\langle f, \gamma(f) \rangle = \sum_{l=0}^{b} \sum_{|m|,|n| \leq l} \left(\sum_{i=1}^{N} a_{l,m}[i] \overline{a_{l,n}[i]} \right) \langle Y_l^m, \gamma(Y_l^n) \rangle.$$

Computing the autocorrelation of f with all of its rotations requires: $< O(Nb^3) = O(b^4)$ time to compute the necessary spherical harmonics, $O(Nb^3) = O(b^4)$ time to compute the cross multiplication of harmonic coefficients, and $\leq O(b^4)$ time to perform the inverse Winger-D transform. Thus, the total complexity of computing the autocorrelation is $O(b^4)$.

Note that increasing the dimensionality of the shape representation from a 2D spherical function to a 3D function does not increase the complexity of computing the symmetry descriptors since the inverse Winger-D transform remains the limiting step.

4.3. Computing the Descriptors

In order to compute the symmetry descriptors of a function f, it suffices to compute the lengths of the projections:

$$\left\| \pi_{G_p^k}(f) \right\|^2 = \frac{1}{|G_p^k|} \sum_{\gamma \in G_p^k} \langle f, \gamma(f) \rangle$$

for all k and all points p on the unit sphere. When k is positive, the elements in G_p^k are all rotations. Thus, having computed the values the autocorrelation, we can reconstruct the symmetry descriptors $\text{SD}_k(f, p)$. However, when k is negative, some elements of G_p^k will be of the form $\gamma = r_p^\alpha \cdot A$ — products of a rotation and the antipodal map. In this case, we cannot use the computed autocorrelation values directly.

In order to be able to compute the symmetry descriptors $\text{SD}_k(f, p)$ for negative k, we observe that the antipodal map acts on a function f as follows:

1. If f is an even function, the antipodal map leaves f unchanged
2. If f is an odd function, the antipodal map sends the function f to the function $-f$.

Thus, we can compute the symmetry descriptors $\text{SD}_k(f, p)$ if we address the even and odd frequencies of f independently. In particular, we express f as the sum of its even and odd components, $f = f^+ + f^-$, with:

$$f^+(p) = \frac{f(p) + f(-p)}{2} \qquad f^-(p) = \frac{f(p) - f(-p)}{2}$$

Then, instead of computing the autocorrelation of f, we compute the autocorrelation of the even and odd parts independently to get:

$$\Phi_f^+(\gamma) = \langle f^+, \gamma(f^+) \rangle \qquad \Phi_f^-(\gamma) = \langle f^-, \gamma(f^-) \rangle$$

This provides a general expression for the value of $\left\| \pi_{G_p^k}(f) \right\|^2$ for all $k \neq 0$ and all axes p as:

$$\frac{1}{|2k|} \left(\sum_{j=1}^{|2k|} \Phi_f^+(r_p^{2j\pi/k}) + (\text{Sgn}(k))^j \Phi_f^-(r_p^{2j\pi/k}). \right)$$

Note that to compute the measure of axial symmetry $\|\pi_{G_p^\infty}(f)\|$, it suffices to compute $\|\pi_{G_p^k}(f)\|$ with k equal to twice the band width of the function.

Complexity: For both spherical functions and 3D functions the complexity of computing the autocorrelation over all rotations is bounded by $O(b^4)$. Since computing the k-th symmetry descriptor requires $O(k)$ summations for each of $O(b^2)$ points on the sphere the overall complexity of computing the $O(b)$ symmetry descriptors is $O(b^4)$.

5. Symmetry and Model Similarity

Work in symmetry detection has been motivated, in part, by the recognition that symmetry is a property characterizing global shape information so that storing a small amount of symmetry information for each model should provide an efficient bound for the similarity of two models. In this section, we formalize this intuition by explicitly describing how the difference in the symmetries of two models relates to their measure of similarity.

5.1. Globality

A fundamental property of the symmetry descriptors is that they characterize global properties of a model, and hence if the symmetry descriptors of two models differ at even one point, we expect this to imply that the models must be different. This property can be formulated explicitly by stating that the L_∞-difference between symmetry descriptors bounds the L_2-difference of the models:

$$\max_k \left\| \text{SD}_k(f,p) - \text{SD}_k(g,p) \right\|_\infty \leq \left\| f-g \right\|_2.$$

The explicit proof of this bound derives from the fact that the values of the symmetry descriptors of a function are equal to the lengths of its projections onto two orthogonal subspaces. Hence, for any k-fold symmetry, and any axis p we have:

$$
\begin{aligned}
\left\| f-g \right\|^2 &= \left\| \pi_{G_p^k}(f) - \pi_{G_p^k}(g) \right\|^2 + \left\| \pi_{G_p^k}^\perp(f) - \pi_{G_p^k}^\perp(g) \right\|^2 \\
&\geq \left(\left\| \pi_{G_p^k}(f) \right\| - \left\| \pi_{G_p^k}(g) \right\| \right)^2 + \\
&\qquad + \left(\left\| \pi_{G_p^k}^\perp(f) \right\| - \left\| \pi_{G_p^k}^\perp(g) \right\| \right)^2
\end{aligned}
$$

$$= \left\| \text{SD}_k(f,p) - \text{SD}_k(g,p) \right\|^2$$

so that the difference between the symmetry descriptors of two models, at any point and any type of symmetry, is an explicit bound for the proximity of the two models.

5.2. Continuous Symmetry Classification

One of the challenges of shape retrieval stems from the fact that often 3D models are not a priori aligned, and many methods for comparing two models require an initial step of pair-wise registration. For these types of applications, the globality property mentioned above cannot be utilized without first aligning the models. In this section we show how symmetry information can be used for comparing two models without requiring the initial alignment step.

We are motivated in our approach by early work in symmetry detection [Ata85, Wol85, Hig86] where the goal was to classify models in terms of the types of symmetry that they have. These methods sought to assign a binary value to each integer k, indicating whether or not a model had k-fold symmetry. Since such a representation did not specify the axis of symmetry, it was inherently rotation invariant.

Using the symmetry descriptors, we extend these binary classifications into a continuous framework where for each k, we store the optimal measure of k-fold symmetry, even when the model is not k-fold symmetric. In particular, setting $s_k(f)$ to be the maximal value of k-fold symmetry of f:

$$s_k(f) = \max_{p \in S^2} \left\| \pi_{G_p^k}(f) \right\|$$

we define the optimal k-fold symmetry of f as the pair:

$$\text{Sym}_k(f) = \left(s_k(f), \sqrt{\|f\|^2 - s_k(f)^2} \right)$$

giving a continuous, rotation invariant classification of a model in terms of its symmetries. Furthermore, as a direct corollary of the globality property, it follows that the symmetry classification can be used to bound the proximity of two models:

$$\max_k \left\| \text{Sym}_k(f) - \text{Sym}_k(g) \right\| \leq \left\| f-g \right\|.$$

Thus, symmetry classifications can be used to match models without requiring an initial step of pair-wise registration.

6. Symmetry Augmentation

Motivated by the property described in Section 5, we would like to use the continuous symmetry classification for efficiently comparing models in a rotation invariant manner. In particular, we would like to augment existing shape descriptors with symmetry information, but would like to do so in a manner that is not redundant. To this end, we consider the Spherical Harmonic Representation described in [Kaz03].

The Spherical Harmonic Representation is a general

method for obtaining a rotation invariant representation of spherical (and 3D) shape descriptors that describes the descriptors in terms of the distribution of energies across different frequencies (and radii). Specifically, given a spherical function f, the Spherical Harmonic Representation decomposes the function in terms of its frequency components f_l:

$$f = \sum_{l=0}^{b} f_l, \qquad \text{with} \qquad f_l = \sum_{|m| \leq l} a_{l,m} Y_l^m$$

where $a_{l,m}$ are the coefficients of f with respect to the spherical harmonic basis Y_l^m. A rotation invariant representation of f is then obtained by storing only the norms of the different frequency components:

$$\text{SH}(f) = \{\|f_0\|, \|f_1\|, \ldots, \|f_b\|\}.$$

The advantages of this representation are two-fold: First, the representation is rotation invariant by construction, making it possible to compare models without first aligning them. Second, in going from a spherical function to its Spherical Harmonic Representation, the dimensionality of the representation is reduced, contracting a 2D spherical function to a 1D array of energy values.

However, it has been noted that the Spherical Harmonic Representation treats each frequency component independently and does not capture information characterizing the alignment between different frequency components. Symmetry, by contrast, depends strongly on the manner in which the different frequencies align, and therefore captures information that is missing in the Spherical Harmonic Representation. Thus, augmenting the Spherical Harmonic Representation with symmetry information should provide a more discriminating representation, combining the local (in frequency space) information of the Spherical Harmonic Representation with global symmetry information.

Figure 2 demonstrates the motivation for this approach. In this figure, a database is queried with the near-axially symmetric table on the left, and retrieval results are shown without (top) and with (bottom) symmetry augmentation. Note that the addition of symmetry induces a preference for models that are near-axially symmetric, and pushes away models (such as the square table, second model in the non-augmented results) that do not have such a symmetry.

In order to augment the Spherical Harmonic Representation we make the assumption that symmetry is uniformly distributed across all the non-constant frequencies (constant frequency components are fully symmetric), so that if f is a shape descriptor and $\text{Sym}_k(f)$ is the measure of the k-fold symmetry of f then:

$$\text{Sym}_k(f_l) \approx \text{Sym}_k(f) \cdot \frac{\|f_l\|}{\|f\|}$$

where f_l is the l-th frequency component of f. Thus, we replace the original Spherical Harmonic Representation (SH)
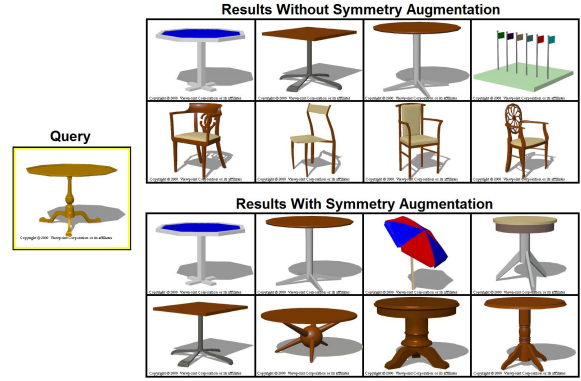

Figure 2: *An example of the improvement gained by augmenting the energy representation with symmetry information. The database was queried with the near axially-symmetric table on the left and results are shown for retrieval without (top) and with (bottom) symmetry augmentation (considering, axial, 2-fold, 3-fold, and reflective symmetries). Note that symmetry augmentation improves matching performance by introducing a preference for models which have near axial symmetry.*

with the symmetry augmented representations:

$$\text{SH}_k(f) = \left\{ \|f_0\|, \text{Sym}_k(f) \cdot \frac{\|f_1\|}{\|f\|}, \ldots, \text{Sym}_k(f) \cdot \frac{\|f_b\|}{\|f\|} \right\}.$$

Then, to compare two descriptors, we find the symmetry type for which the two models vary most, and compare the corresponding symmetry augmented representations:

$$D(f,g) = \max_k \|\text{SH}_k(f) - \text{SH}_k(g)\|.$$

Figure 3 demonstrates the process of symmetry augmentation. Given a spherical shape descriptor (shown in the top left), its Spherical Harmonic Representation is computed by expressing the spherical function in terms of its frequency components, $\{f_0, f_1, \ldots\}$, and storing the norm of each component (shown in the top right). The symmetry descriptors are computed, and the continuous, k-fold symmetry of f is extracted (shown in the bottom left). Finally, the Spherical Harmonic Representation is augmented with symmetry information by scaling with the k-fold symmetry of f, to obtain a finer resolution of non-constant frequency information (shown in bottom right).

6.1. Comparing the Symmetry Augmented Descriptor

Despite the fact that the symmetry augmented representation now requires a copy of the Spherical Harmonic Representation for each symmetry type, in theory encumbering both storage and comparison, the symmetry augmented representation is in fact compact and compares efficiently. In particular, if we compute the symmetry dot product:

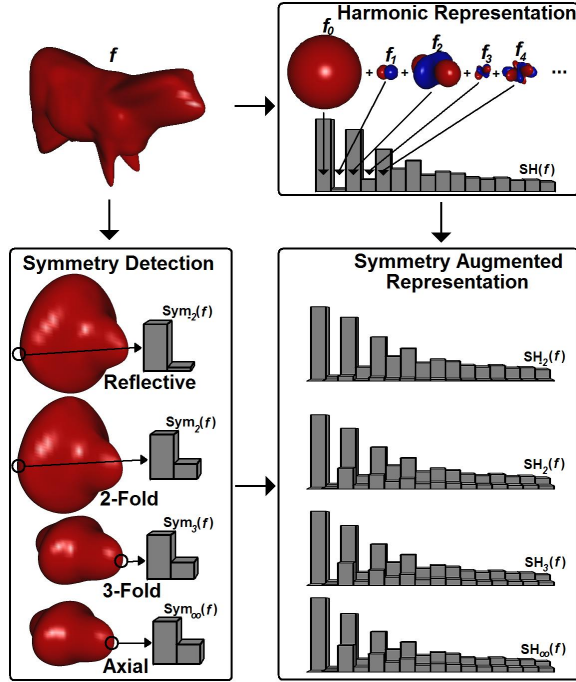$$\text{SDot}(f,g) = \max_k \left\langle \frac{\text{Sym}_k(f)}{\|f\|}, \frac{\text{Sym}_k(g)}{\|g\|} \right\rangle$$

Figure 3: *The augmented Spherical Harmonic Representation of a 3D shape descriptor (top left) is obtained by first computing the Spherical Harmonic Representation (top right) and the k-fold symmetries of f (bottom left). The k-fold symmetries are then used to provide a finer resolution of non-constant frequency information by multiplying each frequency norm by the pair of k-fold symmetry values (bottom right).*

and the frequency dot product:

$$\text{FDot}(f,g) = \sum_{l=1}^{b} \|f_l\| \|g_l\|$$

independently, we can separate the role of symmetry information from frequency information in the measure of shape similarity:

$$D^2(f,g) = \|f\|^2 + \|g\|^2$$
$$-2\|f_0\| \|g_0\| - 2\text{SDot}(f,g) \cdot \text{FDot}(f,g).$$

Thus, in comparing two descriptors, the symmetry information is separated from frequency information and only a single copy of the Spherical Harmonic Representation needs to be stored. Furthermore, the separation of symmetry information from frequency information allows for efficient comparison of two models, since the computations of $\text{SDot}(f,g)$ and $\text{FDot}(f,g)$ are both efficient computations that can be performed independently, and then combined to give the measure of similarity.

Finally, the separation of symmetry information from frequency information provides an easy method for modulating

the importance of symmetry in the measure of model similarity. In particular, we can define the family of metrics:

$$D_{\alpha}^2(f,g) = \|f\|^2 + \|g\|^2$$
$$-2\|f_0\| \|g_0\| - 2\text{SDot}^{\alpha}(f,g) \cdot \text{FDot}(f,g).$$

indexed by the parameter α. When $\alpha = 0$ symmetry plays no role in shape comparison and we revert to the Spherical Harmonic Representation. When $\alpha = 1$ we obtain the symmetry augmented representation described above. And more generally, as α is increased, symmetry plays a more defining role in evaluation of shape similarity.

7. Experimental Results

To measure the efficacy of the symmetry augmented Spherical Harmonic Representation in tasks of shape retrieval, we computed a number of spherical shape descriptors, and compared matching results when the Spherical Harmonic Representation was used with the results obtained when the Spherical Harmonic Representation was augmented with symmetry information. The descriptors we used in our experiments were:

- **Spherical Extent Function**[Vra1]: 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. (If the intersection is empty then the associated value is zero.)
- **Radialized Spherical Extent Function**[Vra03]: A description of a surface which first decomposes 3D space into concentric shells and then computes the Spherical Extent Function of the intersection of the model with each shell independently. The resulting descriptor represents a 3D model with a collection of spherical functions.
- **Shape Histogram (Sectors)**[Ank99]: A description of a surface associating to each ray from the origin, the amount of surface area that sits over it.
- **Shape Histogram (Sectors and Shells)**[Ank99]: A description of a surface which first decomposes 3D space into concentric shells and then computes the Sector representation of the intersection of the model with each shell independently. The resulting descriptor represents a 3D model with a collection of spherical functions.
- **Voxel**: A description of a model as a voxel grid, which is obtained by rasterizing the boundary of the model. The voxel grid is represented as a collection of spherical functions by restricting the grid to concentric spheres and scaling by $\sqrt{4\pi \cdot r^2}$ to account for the change of area.
- **Gaussian Euclidean Distance Transform**[Fun03]: A description of a shape as a voxel grid, where the value at each point is given by the composition of a Gaussian with the Euclidean Distance Transform of the surface. Similar to the *Voxel* representation, the grid is represented by a collection of spherical functions.

We evaluated the performance of each method by measuring how well they classified models within a test database.
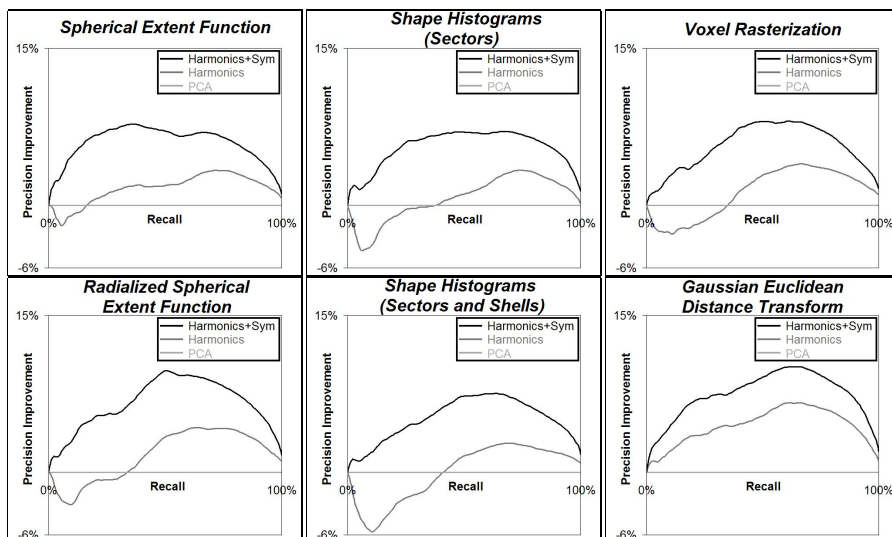
Figure 4: *The improvement in the precision of both the original and augmented Spherical Harmonic Representation over PCA-alignment are shown. Note that symmetry augmentation always improves the matching performance and out-performs the PCA-aligned representation.*

The database consisted of 1890 "household" objects provided by Viewpoint [Vie01]. The objects were clustered into 85 classes, based on functional similarity, 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 [Fun03]. Classification performance was measured using precision/recall plots, which give the percentage of retrieved information that is relevant as a function of the percentage of relevant information retrieved. That is, for each target model in class C and any number K of top matches, "recall" represents the ratio of models in class C returned within the top K matches, while "precision" indicates the ratio of the top K matches that are in class C. Thus, plots that appear shifted up generally indicate superior retrieval results.

The results of the Precision vs. Recall experiments are shown in Figure 4, where the Spherical Harmonic Representation is augmented with k-fold symmetry information, with $k = -2, 2, 3, 4, 5, \infty$ corresponding to reflective, 2-fold, 3-fold, 4-fold, 5-fold and axial symmetry information. A value of $\alpha = 2$ was used to amplify the importance of symmetry in retrieval – this was empirically determined to give the best results. The plots compare the improvement in precision over PCA-alignment when models are matched by comparing the (rotation invariant) Spherical Harmonic Representations of the descriptors and models are matched by comparing the symmetry augmented Spherical Harmonic Representations of the descriptors, (so that PCA-aligned models would give a constant base-line of 0% improvement.) Note that for all of the descriptors, the symmetry augmented representation provides better matching performance, improv-

ing on the retrieval performance of the original Spherical Harmonic Representation.

Representation		Spherical	Voxel
PCA	Size (floats)	256	8192
	Compute (sec)	0.01	0.09
	Compare (ms)	0.18	5.89
Harmonic	Size (floats)	16	512
	Compute (sec)	0.01	0.09
	Compare (ms)	0.02	0.31
Symmetry	Size (floats)	28	524
	Compute (sec)	0.59	0.72
	Compare (ms)	0.02	0.32

Table 1: *A table of the sizes and compute and compare times of the spherical and voxel shape descriptors using PCA-alignment, Spherical Harmonic Representations, and symmetry augmentation.*

Of particular importance is the fact that augmented representation always outperforms the PCA-aligned descriptors (the improved precision plot is always bigger than zero). For a number of descriptors we find that at low recall values the Spherical Harmonic Representation suffers from the information loss inherent in the representation and performs worse than PCA-aligned descriptors. By contrast, the augmented representation always does better than PCA-alignment, despite the fact that the augmented Spherical Harmonic Representation is much smaller than the PCA-aligned representation. The space and time complexity of the different representations are compared in Table 1 which

gives the size and computation time for spherical and voxel based descriptors (voxels are sampled at 32 concentric shells about the origin and every spherical function is represented by its first 16 frequency components). Thus, the augmented representation provides better matching performance with less information, making it particularly well suited for retrieval tasks where compactness, efficiency, and discrimination are imperative.

8. Conclusion

In this paper, we have explored the manner in which symmetry can be used to assist in the tasks of shape matching and shape retrieval. To this end we have introduced the *symmetry descriptors* of a model – a collection of functions giving the measures of k-fold symmetry with respect to all axes passing through a model's center of mass – and have shown how to compute the descriptors efficiently. Using the fact that the symmetry descriptors capture global information, we have shown how the local Spherical Harmonic Representation can be augmented with symmetry information to provide a more discriminating representation for many existing shape descriptors. As a result, we provide a method for obtaining a compact, rotation invariant, representation of shape descriptors that allows for efficient matching of 3D models without requiring a priori model registration.

References

[Ank99] ANKERST M., KASTENMÜLLER G., KRIEGEL H., SEIDL T.: 3D shape histograms for similarity search and classification in spatial databases. In *Advances in Spatial Databases*, (1999), 207–226. 7

[Ata85] ATALLAH M. J.: On symmetry detection. *IEEE Trans. on Computers c-34*, 7 (1985), 663–666. 1, 2, 5

[Att55] ATTNEAVE F.: Symmetry, information, and memory for patterns. *American Journal of Psychology 68* (1955), 209–222. 1

[Fun03] 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. 7, 8

[Hig86] HIGHNAM P.: Optimal algorithms for finding the symmetries of a planar point set. *Information Processing Letters 22* (1986), 219–222. 1, 2, 5

[Kaz02] KAZHDAN M., CHAZELLE B., DOBKIN D., FINKELSTEIN A., FUNKHOUSER T.: A reflective symmetry descriptor. *ECCV* (2002), 642–656. 1, 3

[Kaz03] KAZHDAN M., FUNKHOUSER T., RUSINKIEWICZ S.: Rotation invariant spherical harmonic representation of 3D shape descriptors. *Symposium on Geometry Processing* (2003), 167–175. 5

[Kaz04] KAZHDAN M., CHAZELLE B., DOBKIN D., FUNKHOUSER T., RUSINKIEWICZ S.: A reflective symmetry descriptor for 3D models. *Algorithmica: Special Issue* (2004). 1, 3

[Knu77] KNUTH D., J.H. MORRIS J., PRATT V.: Fast pattern matching in strings. *SIAM Journal of Computing 6*, 2 (1977), 323–350. 2

[Kos03] KOSTELEC P., ROCKMORE D.: *FFTs on the Rotation Group*. Tech. Rep. 03-11-060, Santa Fe Institute's Working Paper Series, 2003. 4

[Mar89] MAROLA G.: On the detection of the axes of symmetry of symmetric and almost symmetric planar images. *IEEE PAMI 11*, 1 (1989), 104–108. 1, 2

[Oma96] O'MARA D., OWENS R.: Measuring bilateral symmetry in digital images. In *TENCON Digital Signal Processing Applications* (1996), IEEE, 151–156. 2

[Pop94] POPE A. R.: *Model-Based Object Recognition: A Survey of Recent Research*. Tech. Rep. TR-94-04, University of British Columbia, 1994. 2

[Ser77] SERRE J.: *Linear Representations of Finite Groups*. Springer-Verlag, New York, 1977. 3

[Sof03] SOFT 1.0: www.cs.dartmouth.edu/~geelong/soft/. 4

[Sun95] SUN C.: Symmetry detection using gradient information. *Pattern Recognition Letters 16* (1995), 987–996. 1, 2

[Sun97] SUN C., SHERRAH J.: 3D symmetry detection using the extended Gaussian image. *IEEE PAMI 19*, 2 (1997), 164–168. 2

[Tan04] TANGELDER J., VELTKAMP R.: A Survey of Content Based 3D Shape Retrieval Methods. In *Shape Modeling International* (2004). 2

[Vet92] VETTER T., POGGIO T., BULTHOFF H.: 3D object recognition: Symmetry and virtual views. In *MIT AI Memo* (1992). 1

[Vie01] VIEWPOINT DATA LABS: www.viewpoint.com. 8

[Vra1] VRANIC D., SAUPE D.: 3D model retrieval with spherical harmonics and moments. *Proceedings of the DAGM* (2001), 392–397. 7

[Vra03] VRANIC D.: An improvement of rotation invariant 3D shape descriptor based on functions on concentric spheres. In *IEEE International Conference on Image Processing* (2003). 7

[Wol85] WOLTER J. D., WOO T. C., VOLZ R. A.: Optimal algorithms for symmetry detection in two and three dimensions. *The Visual Computer 1* (1985), 37–48. 1, 2, 5

[Zab94] ZABRODSKY H., PELEG S., AVNIR D.: Continuous symmetry for shapes. *2nd Intl. Workshop on Visual Form* (1994), 594–613. 1, 2

[Zab95] ZABRODSKY H., PELEG S., AVNIR D.: Symmetry as a continuous feature. *IEEE PAMI 17*, 12 (1995), 1154–1156. 1, 2