

A Reflective Symmetry Descriptor

Michael Kazhdan,
Bernard Chazelle, David Dobkin, Adam Finkelstein, and Thomas Funkhouser

Princeton University, Princeton NJ 08544, USA

Abstract. Computing reflective symmetries of 2D and 3D shapes is a classical problem in computer vision and computational geometry. Most prior work has focused on finding the main axes of symmetry, or determining that none exists. In this paper, we introduce a new *reflective symmetry descriptor* that represents a measure of reflective symmetry for an arbitrary 3D voxel model for all planes through the model’s center of mass (even if they are not planes of symmetry). The main benefits of this new shape descriptor are that it is defined over a canonical parameterization (the sphere) and describes global properties of a 3D shape. Using Fourier methods, our algorithm computes the symmetry descriptor in $O(N^4 \log N)$ time for an $N \times N \times N$ voxel grid, and computes a multiresolution approximation in $O(N^3 \log N)$ time. In our initial experiments, we have found the symmetry descriptor to be useful for registration, matching, and classification of shapes.

1 Introduction

Detecting symmetry in 3D models is a well studied problem with applications in a large number of areas. For instance, the implicit redundancy in symmetric models is used to guide reconstruction [1, 2], axes of symmetry provide a method for defining a coordinate system for models [3], and symmetries are used for shape classification and recognition [4, 5].

Despite its intuitive appeal, symmetry has been under-utilized in computer-aided shape analysis. Most previous methods have focused only on discrete detection of symmetries – i.e., classifying a model in terms of its symmetry groups (either a model has a symmetry, or it does not) [2, 6–13]. Accordingly, they provide limited information about the overall shape of an object, and they are not very useful for shapes that have no symmetries. In contrast, in the context of shape analysis, we believe that it is just as important to know that a model does not have a particular symmetry, as it is to know that it does.

The objective of our work is to define a continuous measure of reflective symmetry (over any plane for any 3D model) and use it to build a concise shape signature that is useful for registration, matching, and classification of 3D objects. Our approach is to define a *reflective symmetry descriptor* as a 2D function that gives the measure of invariance of a model with respect to reflection about each plane through the model’s center of mass. For example, Figure 1 shows a car, cube and chair (top) and their corresponding reflective symmetry descriptors (bottom). The descriptors are drawn by scaling unit vectors on the sphere in proportion to the measure of reflective symmetry about the

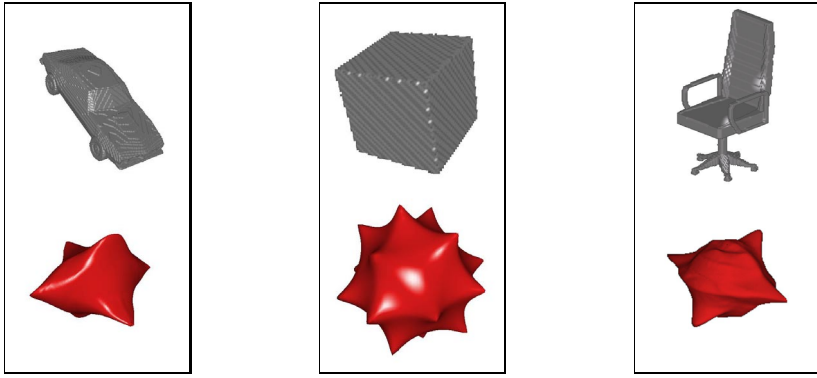


Fig. 1. A visualization of the reflective symmetry descriptor for a car, a cube, and a chair. The visualization is obtained by scaling unit vectors on the sphere in proportion to the measure of reflective symmetry about the plane through the center of mass, normal to the vector.

plane through the center of mass and normal to the vector. Note that the reflective symmetry descriptor provides a continuous measure of reflective symmetry for all planes through the center of mass, and has peaks corresponding to the planes of symmetry, or near symmetry, of the models. For example, the symmetry descriptor of the chair in Figure 1 has strong peaks corresponding to its left-right symmetry, it has smaller horizontal and vertical peaks corresponding to the seat and back of the chair, but it also has a strong peak corresponding to the plane that reflects the back of the chair into the seat. Thus, in the case of the chair, the reflective symmetry descriptor describes not only the different parts of the chair, but also their spatial relationships.

For shape analysis tasks, the potential advantages of the reflective symmetry descriptor are four-fold. First, it characterizes the global shape of the object, and thus it is well-suited for the matching of whole objects (as is often needed for searching large databases of 3D objects). Second, it is defined over a canonical 2D domain (the sphere), and thus it provides a common parameterization for arbitrary 3D models that can be used for alignment and comparison. Third, it is insensitive to noise and other small perturbations in a 3D model, since each point on the symmetry descriptor represents an integration over the entire volume, and thus similar models which only differ in their fine details have similar symmetry descriptors. Finally, it describes the shape of an object in terms of its symmetry features, which provide distinguishing shape information for many objects (look around your office and consider classifying objects based on their symmetries). This approach is quite different from existing shape descriptors, and thus in addition to being useful on its own, it may be helpful to use in conjunction with other representations.

In this paper, we describe our initial research in defining, computing, and using reflective symmetry descriptors. Specifically, we make the following contributions: (1) we define a new continuous measure for the reflective symmetry of a 3D voxel model with respect to a given plane, (2) we describe efficient algorithms to compute the reflective symmetry measure for all planes through the center of mass of a 3D model, and (3) we present experimental results evaluating the utility of reflective symmetry descrip-

tors for registration and classification of 3D models. In our tests, we find that reflective symmetry descriptors are more effective than commonly used shape descriptors (e.g., moments [14] and shape distributions[15]) for registering and classifying 3D models.

The remainder of the paper is organized as follows. Section 2 contains a brief review of related work. Next, Section 3 introduces our new measure of reflective symmetry for voxel models, and Section 4 describes an efficient algorithm for computing a shape descriptor based on this measure. Section 5 discusses some of the properties of the reflective symmetry descriptor, while Section 6 presents experimental results acquired during tests in shape registration and classification applications. Finally, Section 7 contains a brief summary of our work and a discussion of topics for future work.

2 Related Work

Existing approaches for reflective symmetry detection have mainly focused on finding perfect symmetries of a 2D or 3D model [6, 7]. For instance, early work in this area is based on efficient substring matching algorithms (e.g., [16]). However, since substring matching is inherently a binary question, these algorithms can only find perfect symmetries and are highly unstable in the presence of noise and imprecision; thus they are not suitable for most shape registration and matching applications.

In the case of voxel grids, methods for symmetry detection have been proposed using the covariance matrix [11, 12], taking advantage of the fact that eigenspaces of the covariance matrix must be invariant under the symmetries of the model. These methods are efficient and work in all dimensions but have the disadvantage that they only work when the eigenspaces of the covariance matrix are all one-dimensional. In the case of the cube, for example, the covariance matrix is a constant multiple of the identity, every vector is an eigenvector, and no candidate axes of symmetry can be determined. Additionally, the covariance matrix can only identify candidate axes and does not determine a measure of symmetry. So, further evaluation needs to be performed to establish the quality of these candidates as axes of symmetry. Methods for symmetry detection in 2D using more complex moments and Fourier decomposition have also been described [8–10, 13], though their dependence on the ability to represent an image as a function on the complex plane makes them difficult to generalize to three-dimensions.

In the work most similar to ours, Marola [8] presents a method for measuring symmetry invariance of 2D images. However, because of its use of autocorrelation, the method cannot be extended directly to three-dimensional objects. In related work, Zabrodsky, Peleg and Avnir [2] define a continuous *symmetry distance* for point sets in any dimension. Unfortunately, it relies on the ability to first establish point correspondences, which is generally difficult. Additionally, while the method provides a way of computing the symmetry distance for an individual plane of reflection, it does not provide an efficient algorithm for characterizing a shape by its symmetry distances with respect to multiple planes.

Our approach differs from previous work on symmetry detection in that we aim to construct a *shape descriptor* that can be used for registration, matching, and classification of 3D shapes based on their symmetries. The key idea is that the measure of symmetry with respect to any plane is an important feature of an object’s shape, even if

the plane does not correspond to a reflective symmetry of the shape. By capturing this idea in a structure defined on a canonical parameterization, we can compare models by comparing their symmetry descriptors. This basis for comparison provides a means for shape registration, matching, and classification. In this respect, our goals are similar to previously described shape descriptors [17–20], of which some recent examples include spin images [21], harmonic shape images [22], shape contexts [23, 24], and Extended Gaussian Images [25]. Similarly, our descriptor is related to several shape representations that characterize symmetries with respect to *local* axes, such as medial axes [26], shock graphs [27], and skeletons [28, 29]. However, our reflective symmetry descriptor differs from these structures in that it characterizes global symmetry features of a 3D model, and thus it provides shape information orthogonal to these other descriptors.

In the following sections, we describe our methods for computing the reflective symmetry descriptor. There are two main challenges. First, we must describe a new notion of symmetry distance that can be used to measure the invariance of a 3D voxel model with respect to reflection about any plane. Second, we must develop an algorithm for computing the reflective symmetry descriptor that is more efficient than the brute force $O(N^5)$, algorithm for $N \times N \times N$ voxel grids.

3 Defining the Symmetry Distance

The first issue is to define a measure of symmetry for a 3D model with respect to reflection about a plane. While previous work has proposed symmetry measures for 2D images and 3D point sets, we seek such a measure for 3D models based on a solid mathematical framework. This allows us to prove valuable properties of the descriptor.

We define the symmetry distance of a function with respect to a given plane of reflection as the L^2 -distance to the nearest function that is invariant with respect to the reflection. Specifically, we treat a voxel model as a regular sampling of a function and use the L^2 -norm on the space of functions. For a function f and a reflection γ this translates into the equation:

$$\text{SD}(f, \gamma) = \min_{g|\gamma(g)=g} \|f - g\|.$$

Using the facts that the space of functions is an inner product space and that the functions that are invariant to reflection about γ define a vector subspace, it follows that the nearest invariant function g is precisely the projection of f onto the subspace of invariant functions. That is, if we define π_γ to be the projection onto the space of functions invariant under the action of γ and we define π_γ^\perp to be the projection onto the orthogonal subspace then:

$$\text{SD}(f, \gamma) = \|f - \pi_\gamma(f)\| = \|\pi_\gamma^\perp(f)\|$$

so that the symmetry distance of f with respect to γ is the length of the projection of f onto a subspace of functions indexed by γ .

In order to compute an explicit formulation of the projection of f onto the space of functions invariant under the action of γ , we observe that reflections are orthogonal transformations (that is, they preserve the inner product defined on the space of functions). This lets us apply a theorem from representation theory [30] 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 function f and a reflection γ we get:

$$\text{SD}(f, \gamma) = \left\| f - \frac{1}{2}(f + \gamma(f)) \right\| = \left\| \frac{f - \gamma(f)}{2} \right\| \quad (1)$$

so that up to a scale factor the symmetry distance is simply the L^2 -difference between the initial function and its reflection.

As an example, Figure 2 demonstrates this process of projection by averaging. The image on the left shows a picture of Jackie Chan. The image on the right is the closest image that is symmetric with respect to the gray line. It is obtained by averaging the original with its reflection. The L^2 -difference between these two images is the measure of the symmetry of the initial image with respect to reflection about the gray line. Equivalently, according to Equation 1, the symmetry distance is half the L^2 -distance from the original image to its reflection.




Fig. 2. An image of Jackie Chan (left) and its projection onto the space of images invariant under reflection through the gray line (right). The image on the right is obtained by averaging the image on the left with its reflection about the gray line.

4 Computing the Reflective Symmetry Descriptor

The second issue is to build a shape descriptor for a 3D model based on symmetry distances. We define our *reflective symmetry descriptor* as a representation of the symmetry distances for all planes through the model's center of mass. This definition captures global symmetry information and can be parameterized on the domain of a sphere.

In order to use this reflective symmetry descriptor in practical applications, we must develop efficient algorithms for computing it. A naive algorithm would explicitly compute the symmetry distances in $O(N^3)$ time for each of the $O(N^2)$ planes in an $N \times N \times N$ voxel model, requiring $O(N^5)$ time overall. Our approach is to leverage the Fast Fourier Transform to reduce the computation time to $O(N^4 \log N)$ and to use multiresolution methods to provide a good approximation in $O(N^3 \log N)$.

We present our method for efficiently computing the reflective symmetry descriptor of a voxel grid in four steps. First, we show how the Fast Fourier Transform can be used to compute the reflective symmetries of a function defined on a circle efficiently (Section 4.1). Second, we show how the case of a function f defined on the unit disk




Fig. 3. Reflection about α maps a point with angle θ to the point with angle $2\alpha - \theta$.




Fig. 4. The reflective symmetry descriptor of a 2D image can be obtained by decomposing the image into concentric circles and computing the reflective symmetry descriptors on each of the circles.

can be reduced to the case of a function on a circle by decomposing f into a collection of functions defined on concentric circles (Section 4.2). Third, through a collection of mappings we show how to reduce the question of finding the symmetry descriptor of a function on a sphere to the question of finding the symmetry descriptor for a function on a disk (Section 4.3). Fourth, we show how the reflective symmetry descriptor of a voxel grid can be computed by decomposing the grid into a collection of concentric spheres and applying the methods for symmetry detection of functions defined on a sphere (Section 4.4).

We conclude this section by describing a method for efficiently computing a multiresolution approximation to the descriptor, which provides comparable quality at low resolutions in far less time (Section 4.5).

4.1 Functions on a Circle

In order to define the reflective symmetry descriptor for a function on a circle we would like to compute the symmetry distance for reflections about all lines through the origin efficiently. In particular, for a given function f on the circle and any reflection γ we would like to compute the measure of invariance of f with respect to γ . Denoting by γ_α the reflection about the line through the origin with angle α and using the fact that this reflection maps a point with angle θ to the point with angle $2\alpha - \theta$ (see Figure 3) we can apply Equation 1 to obtain:

$$SD(f, \gamma_\alpha) = \sqrt{\frac{L^2\text{-norm}}{2} \left(\underbrace{\|f\|^2}_{\text{convolution term}} - \int_0^{2\pi} \frac{f(\theta)f(2\alpha - \theta)}{2} d\theta \right)}$$

This formulation provides an efficient method for computing the reflective symmetry descriptor of a function defined on a unit circle because we can use the Fast Fourier Transform to compute the value of the convolution term for all angles α in $O(N \log(N))$ time, where N represents the number of times f is sampled on the circle.

4.2 Functions on a Disk

As with functions on a circle, the reflective symmetry descriptor of a function on a disk is a mapping that associates to every angle α the measure of the invariance of the function with respect to the reflection about the line with angle α . To compute the reflective symmetry descriptor we observe that these reflections fix circles of constant radius, and hence the symmetries of a function defined on a disk can be studied by looking at the restriction of the function to concentric circles. Figure 4 shows a visualization of this process where the image of Superman is decomposed into concentric circles and the reflective symmetry descriptor of the image is computed by combining the reflective symmetry descriptors of the different circular functions.

To make this observation explicit we reparameterize the function $f(x, y)$ into polar coordinates to get the collection of functions $\{\tilde{f}_r\}$ with:

$$\tilde{f}_r(\theta) = f(r \cos \theta, r \sin \theta),$$

where $r \in [0, 1]$ and $\theta \in [0, 2\pi]$, and we set γ_α to be the reflection about the line through the origin with angle α . Using Equation 1 and applying the appropriate change of variables we get:

$$\text{SD}(f, \gamma_\alpha) = \sqrt{\int_0^1 \text{SD}^2(\tilde{f}_r, \gamma_\alpha) r dr}$$

showing that we can take advantage of the efficient method for computing the reflective symmetry descriptor of a function on the circle to obtain an $O(N^2 \log N)$ algorithm for computing the reflective symmetry descriptor of an $N \times N$ image.

This method is similar to the method presented in the works of Marola and Sun et al. [8, 10] in its use of autocorrelation as a tool for reflective symmetry detection. The advantage of our formulation is that it describes the relationship between autocorrelation and an explicit notion of symmetry distance, defined by the L^2 inner-product of the underlying space of functions, and provides a method for generalizing the definition of symmetry distance to 3D.

4.3 Functions on a Sphere

The reflective symmetry descriptor of a function defined on the surface of a sphere is a mapping that gives the measure of reflective invariance of a model with respect to reflections about every plane through the origin. To compute the reflective symmetry descriptor of a function on a sphere we fix a North pole and restrict our attention to those planes passing through it. The values of the reflective symmetry descriptor for the restricted set of reflections can be efficiently computed by breaking up the function into its restrictions to the upper and lower hemisphere and projecting each of these restrictions to a disk. Figure 5(left) shows a visualization of this process for the restriction to the upper hemisphere. Note that reflections through planes containing the North pole map the upper hemisphere to itself and correspond to reflections about lines in the projected function.

In particular if we parameterize the sphere in terms of spherical coordinates:

$$\Phi(\phi, \theta) = (\cos \phi, \sin \phi \cos \theta, \sin \phi \sin \theta)$$




Fig. 5. To compute the reflective symmetry descriptor of a function defined on the sphere for planes passing through the North pole we observe that these planar reflections correspond to reflections about lines in the projected functions (left). We observe that a great circle must intersect every plane through the origin (right) so letting the North pole vary over a great circle and computing the projection at every step we obtain the symmetry distance for all planes.

with $\phi \in [0, \pi]$ and $\theta \in [0, 2\pi]$, the restriction to the upper hemisphere corresponds to the restriction $\phi \in [0, \pi/2]$. Unfolding the restriction of f to the upper hemisphere along lines of constant latitude gives a function \tilde{f}_u defined on a disk of radius $\pi/2$:

$$\tilde{f}_u(\phi \cos \theta, \phi \sin \theta) = f(\Phi(\phi, \theta)) \sqrt{\frac{\sin \phi}{\phi}}.$$

We can obtain \tilde{f}_l , the projection of the lower hemisphere, in a similar fashion. Letting γ_α represent both the reflection of the sphere about the plane through the North pole with constant angle of longitude α and the reflection of the disk about the line with angle α we get:

$$\text{SD}(f, \gamma_\alpha) = \sqrt{\text{SD}^2(\tilde{f}_u, \gamma_\alpha) + \text{SD}^2(\tilde{f}_l, \gamma_\alpha)}$$

so that with the correct parameterization and scaling of the projections, the symmetry distance for the reflection γ_α can be obtained from the symmetry distances of the projections of f . (Note that rather than doing a true projection onto the plane perpendicular to the North pole, we actually unfold the hemisphere in terms of its angles of latitude. This allows us to avoid the sampling problems that would otherwise result due to a vanishing Jacobian near the boundary of the disk.)

In order to compute the value of the reflective symmetry descriptor for all planes through the origin, not just those passing through the North pole, we use the fact that if we fix a great circle on the sphere, any plane through the origin must intersect the great circle in at least two points (Figure 5(right)). This allows us to compute the values of the reflective symmetry descriptor for all planes by walking (half of) the great circle and at each point projecting onto a disk to compute the measure of symmetries for those planes containing the current North pole. Since the symmetry descriptor of the projection onto a disk can be computed in $O(N^2 \log N)$ and since we perform $O(N)$ such projections, this method gives an $O(N^3 \log N)$ algorithm for computing the reflective symmetry descriptor of a function on the sphere, sampled at $O(N^2)$ points.

4.4 Functions on a Voxel Grid

As with a function defined on a sphere, the reflective symmetry descriptor of a voxel model is a function that gives the measure of invariance of the model with respect to reflection about every plane through the origin, where we assume that the model's center of mass has been translated to the origin. As in Section 4.2, we can use the fact that reflections fix lengths to transform the problem of computing the reflective symmetry descriptor of a voxel grid into a problem of computing the reflective symmetry descriptors of a collection of functions defined on a sphere. In particular, if f is a function defined on the set of points with radius less than or equal to 1 then we can decompose f into a collection of functions $\{\tilde{f}_r\}$ where \tilde{f}_r is a function defined on the unit sphere and $\tilde{f}_r(v) = f(rv)$. After changing variables, the measure of symmetry of f with respect to a reflection γ becomes:

$$\text{SD}(f, \gamma) = \sqrt{\int_0^1 \text{SD}^2(\tilde{f}_r, \gamma) r^2 dr}$$

and we obtain the value of the symmetry descriptor of f as a combination of the values of the symmetry descriptors of the spherical functions $\{\tilde{f}_r\}$, giving a method for computing the reflective symmetry descriptor of an $N \times N \times N$ model in $O(N^4 \log N)$.

4.5 Multiresolution Approximation

Our algorithm for computing the reflective symmetry descriptor takes $O(N^4 \log N)$ time at full resolution. However, using Fourier decomposition of the restriction of the function to lines through the origin, we are able to compute a good multiresolution approximation to it in $O(N^3 \log N)$ time. This approximation is useful in most applications because symmetry describes global features of a model and is apparent even at low resolutions. Given a function f defined on the set of points with radius less than or equal to 1 we decompose f into the collection of one-dimensional functions by fixing rays through the origin and considering the restriction of f to these rays. This gives a collection of functions $\{\tilde{f}_v\}$, indexed by unit vectors v , with $\tilde{f}_v(t) = f(tv)t$ and $t \in [0, 1]$. Expanding the functions \tilde{f}_v in terms of their trigonometric series we get:

$$\tilde{f}_v(t) = a_0(v) + \sum_{k=1}^{\infty} \left(a_k(v) \frac{\cos(2k\pi t)}{\sqrt{2}} + b_k(v) \frac{\sin(2k\pi t)}{\sqrt{2}} \right).$$

The advantage of this decomposition is that the functions $a_k(v)$ and $b_k(v)$ are functions defined on the sphere, providing a multiresolution description of the initial function f . Applying the appropriate change of variables and letting γ denote a reflection about a plane through the origin we get:

$$\text{SD}(f, \gamma) = \sqrt{\text{SD}^2(a_0, \gamma) + \sum_{k=1}^{\infty} (\text{SD}^2(a_k, \gamma) + \text{SD}^2(b_k, \gamma))}.$$

Thus a lower bound approximation to the reflective symmetry descriptor can be obtained in $O(N^3 \log N)$ time by only using the first few of the functions a_k and b_k .




Fig. 6. (a) A 1957 Chevrolet model, (b-d) the approximations of its symmetry descriptor using the first 4, 8, and 32 spherical coefficient functions, and (e) the descriptor at full resolution.

The advantage of this multiresolution decomposition is that for binary voxel models we can show that the approximations converge quickly to the true value of the reflective symmetry descriptor at every point. In particular, if c represents the complexity of the model (i.e. the number of times a line through the origin will enter and exit the shape) then the approximation using only the first k Fourier coefficient functions differs from the true reflective symmetry descriptor in proportion to c/k . (This follows from the fact that if $g(t)$ is the characteristic function of c disjoint segments contained on the interval $[0, 2\pi]$ then $\int \sin(2k\pi t)g(t)dt \leq c/k$.) This result is demonstrated empirically in Figure 6(b-e), which shows the symmetry descriptor computed for a Chevrolet using 4, 8, 32, and 128 Fourier coefficient functions, respectively. Note that using only the first eight Fourier coefficient functions results in an approximation that is barely distinguishable from the higher resolution versions.

5 Additional Properties of the Reflective Symmetry Descriptor

In addition to being a function that is both parameterized over a canonical domain and describes a model in terms of its symmetries, the reflective symmetry descriptor has provable properties valuable for shape analysis:

Stability: The reflective symmetry descriptor is stable in the presence of high-frequency noise. To see this, we rewrite the reflective symmetry descriptor of a function f , defined on a circle, in terms of its Fourier coefficients:

$$\text{SD}(f, \gamma_\alpha) = \sqrt{\frac{1}{2} \sum_k (||a_k||^2 + a_k^2 e^{i2k\alpha})}.$$

This equation demonstrates that the contribution of different frequencies to the reflective symmetry descriptor depends only on their Fourier coefficients. In contrast, shape descriptors that involve computation of model derivatives, either as normals or gradients [25, 10], have the property of amplifying the contribution of high-frequency components, making them unstable in the presence of high-frequency noise.

Globality: The differences in the reflective symmetry descriptors of two different models at a single point provides a lower bound for the overall similarity of the two models.

The proof of this bound derives from the fact that the symmetry distance of a function f with respect to a reflection γ is defined as the length of the projection $\pi_\gamma^\perp(f)$ (Section 3). Since we know that for any orthogonal projection π and any vectors v and w we have $\|v\| \geq \|\pi(v)\|$ and $\|v - w\| \geq \|\|v\| - \|w\|\|$ it follows that:

$$\|f - g\| \geq \left| \|\pi_\gamma^\perp(f)\| - \|\pi_\gamma^\perp(g)\| \right| = |\text{SD}(f, \gamma) - \text{SD}(g, \gamma)|$$

so that the difference in the values of two symmetry descriptors at a single point provides a lower bound for the L^2 -difference of the corresponding models.

6 Results

In this section, we show the reflective symmetry descriptors of a wide variety of models and demonstrate the efficacy of the descriptor as a shape analysis tool by showing how models can be registered and classified using the descriptor.

6.1 Test Database

Our test database consisted of 90 polygonal models categorized by a third party into 24 different classes. Using a simple rasterization method, the interior of each model was voxelized into a $128 \times 128 \times 128$ grid.

Figures 7 and 8 show a number of models from the test database with their corresponding reflective symmetry descriptors. The full symmetry descriptor for each model was computed in 122 seconds on an 800 MHz Athlon processor with 512 MB of RAM. Note that the descriptors vary from model to model, with different patterns of undulations and sharp peaks, demonstrating that the symmetry descriptor is a rich function, capable of describing large amounts of information about shape.




Fig. 7. A number of models from different classes with their reflective symmetry descriptors, demonstrating the variability and richness of the descriptor.




Fig. 8. A number of representative models with their corresponding symmetry descriptors. Note that the descriptor remains consistent within a class.

6.2 Registration Results

In this section, we test how well the symmetry descriptors of models can be used for registration. To do this we hand-aligned all pairs of models within a class. For each pair of models, we computed the axis of the rotation which brings the two models into alignment and then searched for the rotation along that axis that minimizes the L^2 -distance between the corresponding symmetry descriptors.

Figure 9 compares errors in registration using the reflective symmetry descriptor with errors using principal axis alignment. The graphs show the percentage of pairwise registrations that resulted in a given error, where error is measured as the difference in the rotation angles between the user-specified and the computed rotations. Note that reflective symmetry descriptors register different models within the same class to within 5 degrees of what a human would do in 95% of the tests as opposed to the covariance approach that only registers to within 5 degrees 55% of the time.

These results indicate that registration using the reflective symmetry descriptor does a better job of aligning models than the classic principal axis method. We believe that the trouble with using principal axis alignment is two-fold. First, in the case that an eigenspace of the covariance matrix is more than one-dimensional, a unique eigenvector cannot be determined. Second, the contribution of points to the covariance matrix scales quadratically with their distance from the center, so that small changes in a model that occur far from the center can drastically change the principal axes. Figure 10 demonstrates this by showing the principal axes of three different mugs. Note that changes in the position and shape of the handle, and changes in the cylindrical nature of the mug give rise to principal axes that are differently aligned. The reflective symmetry descriptor, by contrast, remains stable throughout these variances (Figure 8 bottom row.)




Fig. 9. A comparison of registration using the reflective symmetry descriptor (left) with registration using principal axes (right). The graphs show the percentage of pairwise registration that resulted in a given rotation error.




Fig. 10. A collection of mugs with their principal axes. The figure demonstrates that minor variances within a class can drastically affect the orientation of the principal axes.

6.3 Classification Results

We also evaluated the discriminating power of the reflective symmetry descriptor with respect to the task of object classification. In order to do this efficiently, we generated rotation-invariant signatures for the symmetry descriptor based on the distribution of their values. In particular, we obtained the first eight approximating Fourier coefficient functions (as describes in Section 4.5) and generated histograms for each of their reflective symmetry descriptors. The histograms contained 100 bins, with the k -th bin containing the measure of the points on the sphere whose corresponding symmetry distance was in the range $[k/100, (k + 1)/100]$. We measured model similarity by comparing the obtained histograms using the Earth Mover’s Distance [31].

We performed a sequence of leave-one-out classification experiments for each model based on the measure of model similarity. Table 1 compares the results of model classification using the symmetry distribution with the classification results obtained using higher order moments [14] and shape distributions [15], two other global shape descriptors used for matching and classification. In order to provide a base measure of performance, the table also presents the results when the similarity measure returns a random number. The quality of the classification was measured using three metrics [15]. The *Nearest Neighbor* value is the percentage of models whose closest match belonged to the same class. The *First Tier* and *First Two Tiers* values corresponds to the percent-

age of models in the first $(n - 1)$ and $2(n - 1)$ nearest matches that belonged to the same class as the query model, where n is the class size. They provide measures of classification that normalize for the number of models within a class. Note that reflective symmetry descriptors classify models as well or better than the other shape descriptors for all three classification criteria.

Table 1. Comparison of results of the model matching experiment using 4th and 7th order moments, shape distributions and symmetry descriptors.

Comparison Method	Nearest Neighbor	First Tier	First Two Tiers	Time
Random	6%	4%	9%	
Moments (4th Order)	34%	40%	48%	0.1 seconds
Moments (7th Order)	24%	33%	38%	0.25 seconds
Shape Distributions	44%	64%	62%	0.35 seconds
Symmetry Descriptors	52%	69%	71%	.15 seconds

7 Conclusion and Future Work

In this paper, we have introduced the *reflective symmetry descriptor*, a function associating a measure of reflective invariance of a 3D model with respect to every plane through the center of mass. It has several desirable properties, including invariance to translation and scale, parameterization over a canonical domain, stability, and globality that make it useful for registration and classification of 3D models. We have shown how to compute it efficiently, and conducted preliminary experiments that show its usefulness for shape registration and classification.

This work suggests a number of questions that we would like to address in future research: (1) Can the symmetry descriptor be used for other shape analysis tasks, such as learning a statistical classifier of shape? (2) Can the multiresolution properties of the descriptor be used to develop more efficient search algorithms, e.g., for registration and recognition? (3) Can other theoretical properties of the descriptor be proven, such as showing when 3D models can have the same descriptor? Answers to these questions will further our understanding of how symmetry defines shape.

References

1. Mitsumoto, H., Tamura, S., Okazaki, K., Kajimi, N., Fukui, Y.: Reconstruction using mirror images based on a plane symmetry recovery method (1992)
2. Zabrodsky, H., Peleg, S., Avnir, D.: Symmetry as a continuous feature. *IEEE PAMI* **17** (1995) 1154–1156
3. Liu, Y., Rothfus, W., Kanade, T.: Content-based 3d neuroradiologic image retrieval: Preliminary results (1998)
4. Leou, J., Tsai, W.: Automatic rotational symmetry determination for shape analysis. *Pattern Recognition* **20** (1987) 571–582

5. Wolfson, H., Reisfeld, D., Yeshurun, Y.: Robust facial feature detection using symmetry. *Proceedings of the International Conference on Pattern Recognition* (1992) 117–120
6. Atallah, M.J.: On symmetry detection. *IEEE Trans. on Computers* **c-34** (1985) 663–666
7. 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
8. Marola, G.: On the detection of the axes of symmetry of symmetric and almost symmetric planar images. *IEEE PAMI* **11** (1989) 104–108
9. Shen, D., Ip, H., Cheung, K., Teoh, E.: Symmetry detection by generalized complex (gc) moments: A close-form solution. *IEEE PAMI* **21** (1999) 466–476
10. Sun, C., Si, D.: Fast reflectional symmetry detection using orientation histograms. *Real-Time Imaging* **5** (1999) 63–74
11. O'Mara, D., Owens, R.: Measuring bilateral symmetry in digital images. *IEEE-TENCON - Digital Signal Processing Applications* (1996)
12. Sun, C., Sherrah, J.: 3-d symmetry detection using the extended Gaussian image. *IEEE PAMI* **19** (1997)
13. Kovese, P.: Symmetry and asymmetry from local phase. *Tenth Australian Joint Conference on Artificial Intelligence* (1997) 2–4
14. Elad, M., Tal, A., Ar, S.: Directed search in a 3d objects database using svm (2000)
15. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Matching 3d models with shape distributions. *Shape Matching International* (2001)
16. Knuth, D., J.H. Morris, J., Pratt, V.: Fast pattern matching in strings. *SIAM Journal of Computing* **6** (1977) 323–350
17. Besl, P.J., Jain, R.C.: Three-dimensional object recognition. *Computing Surveys* **17** (1985) 75–145
18. Loncaric, S.: A survey of shape analysis techniques. *Pattern Recognition* **31** (1998) 983–1001
19. Pope, A.R.: Model-based object recognition: A survey of recent research. *Technical Report TR-94-04, University of British Columbia* (1994)
20. Veltkamp, R.C., Hagedoorn, M.: State-of-the-art in shape matching. *Technical Report UU-CS-1999-27, Utrecht University, the Netherlands* (1999)
21. Johnson, A., Hebert, M.: Efficient multiple model recognition in cluttered 3-d scenes. *IEEE CVPR* (1998) 671–677
22. Zhang, D., Hebert, M.: Harmonic maps and their applications in surface matching. *IEEE CVPR* **2** (1999)
23. Belongie, S., Malik, J.: Matching with shape contexts. *IEEE Workshop on Content-based access of Image and Video-Libraries* (2000)
24. Mori, G., Belongie, S., Malik, H.: Shape contexts enable efficient retrieval of similar shapes. *CVPR* **1** (2001) 723–730
25. B. Horn, B.: Extended gaussian images. *PIEEE* **72** (1984) 1656–1678
26. Fernandez-Vidal, S., Bardinet, E., Malandain, G., Damas, S., de la Blanca Capilla, N.: Object representation and comparison inferred from its medial axis. *ICPR* **1** (2000) 712–715
27. Siddiqi, K., Shokoufandeh, A., Dickinson, S., Zucker, S.: Shock graphs and shape matching. *IJCV* **35** (1999) 13–32
28. Bloomenthal, J., Lim, C.: Skeletal methods of shape manipulation. *Shape Modeling and Applications* (1999) 44–47
29. Storti, D., Turkiyyah, G., Ganter, M., Lim, C., Stal, D.: Skeleton-based modeling operations on solids. *Symposium on Solid Modeling and Applications* (1997) 141–154
30. Serre, J.: *Linear Representations of Finite Groups*. Springer-Verlag, New York (1977)
31. Rubner, Y., Tomasi, C., Guibas, L.: A metric for distributions with applications to image databases. *IEEE ICCV* (1998)