

Three-dimensional bilateral symmetry plane estimation in the phase domain

Ramakrishna Kakarala, Prabhu Kaliamoorthi, and Vittal Premachandran
School of Computer Engineering, Nanyang Technological University, Singapore

ramakrishna@ntu.edu.sg, prab0009@e.ntu.edu.sg, vitt0002@e.ntu.edu.sg

Abstract

We show that bilateral symmetry plane estimation for three-dimensional (3-D) shapes may be carried out accurately, and efficiently, in the spherical harmonic domain. Our methods are valuable for applications where spherical harmonic expansion is already employed, such as 3-D shape registration, morphometry, and retrieval. We show that the presence of bilateral symmetry in the 3-D shape is equivalent to a linear phase structure in the corresponding spherical harmonic coefficients, and provide algorithms for estimating the orientation of the symmetry plane. The benefit of using spherical harmonic phase is that symmetry estimation reduces to matching a compact set of descriptors, without the need to solve a correspondence problem. Our methods work on point clouds as well as large-scale mesh models of 3-D shapes.

1. Introduction

The motivation to apply spherical harmonic expansion to solve three-dimensional (3-D) computer vision problems stems from at least three well-known properties. First, the expansion summarizes a large number of shape points (vertices or surface voxels) in a relatively small number of coefficients. Second, the expansion allows separation of coarse and fine levels of detail in the shape through bandwidth. Third, the spherical harmonic coefficients behave predictably under 3-D rotation. Accordingly, spherical harmonics have been successfully employed in computer vision for 3-D shape registration [10], morphometry [3], and recognition [7][9]. In this paper, we show that spherical harmonics also provide an accurate and efficient solution for estimating the bilateral symmetry plane of a 3-D shape. Bilateral symmetry has been studied in numerous works (for example, [12] [6][20]). However, the problem of estimating the symmetry plane through spherical harmonic coefficients alone has not been solved, though there are results for the special case of moment coefficients [11]. Consequently, with current techniques, if the spherical harmonic expansion has already been computed, as would be the case in




Figure 1: The image in (a) is reconstructed from only its Fourier phase in (b), illustrating phase’s importance to appearance. The same is true for the sphere: the continental edges in the world map (c) are clearly visible in (d), which is reconstructed from (c) using only the spherical harmonic “phase” as defined in this paper.

the applications mentioned above, there is no way to reuse the computationally-expensive expansion to obtain a symmetry plane estimate. And yet, as we show, the information required to obtain the estimate is clearly available in the harmonic coefficients, and a relatively simple estimation algorithm is possible.

Our approach makes use of the “phase” of spherical harmonics, which is a term without a widely-accepted definition for spherical harmonics. In contrast, phase is well-defined for the ordinary Fourier transform: on the real line, the phase of transform F is ϕ in the polar decomposition $F = |F|e^{j\phi}$. The analogous definition for spherical harmonics must, as we argue below, consider vectors of coefficients and treat the unit vector direction (vector divided by its length) as the phase. It is well-known that for the ordinary Fourier transform, phase defines the locations of edges and is therefore more important for appearance than magnitude. Figure 1 shows that this is true for phase of spherical

harmonic coefficients as well. The figure motivates us to examine what must be true for spherical harmonic phase if symmetry exists in the spatial domain.

Our main contribution shows that bilateral symmetry in spherical data manifests itself as linear phase structure in the spherical harmonic coefficients. Our results allow symmetry to be determined from any type of spherical harmonic expansion, which is valuable since there are several different expansions in use in the literature. We also propose new algorithms for estimating the orientation of the symmetry plane by maximizing the fit of a linear phase equivalent to the harmonic coefficients. In order to improve the fit, we show how to sample from the surface of a mesh model with a nearly uniform distribution of points. Our methods work for a wide variety of data, including point clouds and polygonal meshes (open as well as watertight), and they do not require that the meshes be aligned to the symmetry or have a star-shaped property.

2. Previous work

Given a point cloud in 3-D, one method for finding candidates for the bilateral symmetry plane is to evaluate the distance of corresponding points when reflected across each candidate plane. The simplest method of finding such candidate planes relies on the well-known fact that reflection symmetry, if it exists, must preserve the eigenspaces of the point cloud’s covariance matrix. If the 3×3 covariance matrix has distinct eigenvalues, then its three eigenvectors determine orthogonal candidates for the symmetry plane. The covariance method has been used to provide starting points in a limited search for peaks in the correlation of the orientation histogram [20]. However, the covariance does not provide a means of establishing which candidate is the correct one, and is not reliable as a source of candidates if the eigenvalues are nearly equal.

Kazhdan *et al.* [6] describe a reflective symmetry descriptor that is constructed by measuring the norm of the projection of a voxel set onto the space of bilaterally-symmetric sets. The descriptor is computed at each candidate plane’s normal vector by using Fourier series expansion, and therefore constitutes a function defined on the sphere. The shape of the descriptor function agrees visually with the perceived symmetries of 3-D shapes, and is valuable for shape registration and classification. However, the problem of estimating the orientation of the symmetry plane is not discussed in [6].

A large number of papers exist on detecting 3-D symmetries, of which [13] provides a survey. The RANSAC algorithm is applied to detect 3-D mirror symmetric objects from images [19]. Symmetry for non-rigid shapes is analyzed and a numerical method for finding corresponding components is provided [17]. The 3-D Fourier transform is calculated on a pseudo-polar grid to detect symmetry

groups in voxel data [1].

The richness of the literature on 3-D symmetry shows that the topic is of considerable interest. While many facets of symmetry have been explored, there is no existing work on determining bilateral symmetry from spherical harmonic expansion alone. However, Martinet *et al.* [11] show that local minima for the spherical harmonic coefficients of even-order moments provide candidates for the symmetry axis. To evaluate the suitability of a candidate reflection R , they compute the maximum possible distance between all vertices of S and their nearest points (which are found by minimization) in the reflection of S by R . That calculation is expensive, and does not take advantage of the structure present in the spherical harmonic coefficients. In this paper, we compute a much cheaper and more suitable distance using spherical harmonics coefficients alone, which is a substantial savings since a shape having thousands of vertices may be represented using only a few tens of coefficients. Moreover, the wide-spread use of spherical harmonics in various applications, as mentioned above, makes it attractive to investigate how to estimate the symmetry plane from harmonic coefficients alone. Our methods are compatible with each of the numerous uses of spherical harmonics in the literature, whether applied to the spherical mappings in [6], to the even-order moments in [11], or to the Zernike coefficient mapping in [9].

3. Notation and background

Conventions and notations for spherical harmonics vary across many papers on 3-D vision, and therefore we specify ours in this section, referring to standard works [21] for details.

We use the z - y - z system of Euler angles, denoted α , β , and γ , to describe each 3-D rotation. The two angles, α and β , also parameterize the sphere S^2 , with $0 \leq \alpha < 2\pi$ as longitude and $0 \leq \beta \leq \pi$ as colatitude: specifically, the 3×1 unit vector X lying on the sphere is

$$X = [\cos(\alpha) \sin(\beta), \sin(\alpha) \sin(\beta), \cos(\beta)]^\top. \quad (1)$$

Spherical harmonics form an orthogonal basis for functions on S^2 , and, for each non-negative integer ℓ , they have the functional form

$$Y_\ell^n(X) = Y_\ell^n(\beta, \alpha) = P_\ell^n(\cos \beta) e^{-jn\alpha}, \quad -\ell \leq n \leq \ell. \quad (2)$$

Here P_ℓ^n are the real-valued *associated Legendre functions*. The expansion of each L_2 function f on the sphere is

$$f(X) = f(\beta, \alpha) = \sum_{\ell=0}^{\infty} \sum_{n=-\ell}^{\ell} F_\ell^n Y_\ell^n(\beta, \alpha) \quad (3)$$

For real-valued functions f , we have conjugate-symmetry in the coefficients: $F_\ell^{-n} = (-1)^n F_\ell^{n*}$. We combine all

coefficients for a given ℓ into a $1 \times (2\ell + 1)$ row vector

$$\mathcal{F}_\ell = [F_\ell^{-\ell}, \dots, F_\ell^0, \dots, F_\ell^\ell]. \quad (4)$$

The vector convention simplifies analysis by removing unnecessary summations. We also write $\mathcal{Y}_\ell(\alpha, \beta)$ to denote the row vector arranged similarly to (4) whose entries are $Y_\ell^n(\alpha, \beta)$ for $n = -\ell, \dots, \ell$.

The rotation property of the spherical harmonic coefficients is important to our development. Every 3-D rotation may be represented by 3×3 orthogonal matrix P with determinant $+1$. Rotation in the spherical harmonic domain is determined for each frequency ℓ by the Wigner matrices $D_\ell(P)$, which are $(2\ell + 1)$ -dimensional, unitary ($D_\ell(P)D_\ell(P)^\dagger = I$ for conjugate-transpose \dagger), and satisfy $D_\ell(PS) = D_\ell(P)D_\ell(S)$ for every pair of rotations P, S . The Wigner matrices play the role corresponding to the complex exponential $e^{j\omega}$ at frequency ω in the ordinary Fourier transform on \mathbb{R} . Specifically, if for some P , it is the case that $g(X) = f(PX)$ for all $X \in S^2$, so that g is a rotated version of f , then the rotation property states that:

$$\mathcal{G}_\ell = \mathcal{F}_\ell D_\ell(P). \quad (5)$$

The elements of the Wigner D -matrices are separable in the Euler angles:

$$D_\ell^{mn}(\alpha, \beta, \gamma) = e^{-jm\alpha} d_\ell^{mn}(\beta) e^{-jn\gamma}, \quad -\ell \leq m, n \leq \ell. \quad (6)$$

Here d_ℓ^{mn} is the (little) Wigner d -function, which is real-valued for the z - y - z choice of Euler angles. Therefore, $D_\ell(0, \beta, 0)$ is a real-valued matrix.

Data in 3-D vision tasks are frequently presented as a set of points (x_i, y_i, z_i) , or equivalently in spherical coordinates $(\alpha_i, \beta_i, \rho_i)$, for $i = 1, \dots, N$. Such data may be approximated with spherical harmonics by choosing coefficients \mathcal{F} to minimize the squared error

$$\sum_{i=1}^N \left\| \rho_i - \sum_{\ell=0}^L \mathcal{F}_\ell \mathcal{Y}_\ell(\alpha_i, \beta_i)^\top \right\|^2. \quad (7)$$

Using the method of Chung *et al.* [3] to obtain the coefficients, the least-squares solution to (7) is illustrated in Figure 2 for two important cases. The first case minimizes (7) by fitting the vertices from a triangulated mesh model, and is referred to as the vertex mapping in this paper. The second case uses the extended Gaussian image (EGI), which associates to each of the triangular faces a surface normal whose length is equal to the area of the face. Other possibilities, which are not shown here but are nevertheless equally applicable, include even-order moment functions [11] and Zernike coefficient mapping [9]. For each case, the spherical harmonic approximation to the mapping function is necessarily a star-shaped surface, but is still able to capture a variety of surface features, as such as the antennae as illustrated for the vertex mapping in part (c) of Fig. 2(c).




Figure 2: The ant model in (a) is approximated in its vertex mapping using two bandwidths in (b) and (c). Note how the features of the ant, especially the antennae, become evident as the top red bumps in (c). The EGI mapping of the model is obtained from surface normals, shown superimposed on the model in (d), and approximated in (e) and (f).

4. Symmetry is equivalent to linear phase

Taking inspiration from the magnitude-phase decomposition of the real-line Fourier transform value as $F = |F|e^{j\phi}$, we define the magnitude-phase decomposition for the spherical harmonic vector \mathcal{F}_ℓ as $\mathcal{F}_\ell = \|\mathcal{F}_\ell\|U_\ell$. The $1 \times (2\ell + 1)$ unit-vector U_ℓ determines the direction of the ℓ -th coefficient vector. From (5), we see that the magnitude $\|\mathcal{F}_\ell\|$ remains invariant under rotation (since the D matrix is unitary), but the spherical harmonic phase rotates as $U_\ell \mapsto U_\ell D_\ell$. Figure 1(d) shows that phase contains significant information about the shape's structure.

The spherical harmonic phase also transforms in a simple way under reflection. Suppose that the reflection is across the x - z axis. Under that reflection, for the (α, β) angular parameterization of the sphere, $\alpha \mapsto -\alpha$. Consequently, (2) shows that the spherical harmonic basis takes complex-conjugate values, and therefore the coefficient vector $\mathcal{F}_\ell \mapsto \mathcal{F}_\ell^*$. Now suppose that the reflection is across a plane with the 3×1 normal vector N ; the plane is denoted N^\perp . Let P be any rotation that maps N to the positive y , *i.e.*, $PN = Y$, where $Y = [0, 1, 0]^\top$ is the y -axis. Using P , we may decompose the reflection across N^\perp into a cascade of three steps: (1) rotation of N to Y by applying P ; (2)




Figure 3: The world map, whose spherical harmonic approximation using $L = 64$ is shown in (a), is reconstructed in (b) from the real part of its coefficients. As expected, the result is bilaterally symmetric.

reflection across x - z plane, *i.e.*, $\alpha \mapsto -\alpha$; (3) rotating the y axis back to N by applying P^\top . Using (5), the cascade is

$$\mathcal{F}_\ell \mapsto \mathcal{F}_\ell D_\ell(P) \mapsto \mathcal{F}_\ell^* D_\ell(P)^* \mapsto \mathcal{F}_\ell^* D_\ell(P)^* D_\ell(P^\top). \quad (8)$$

Consequently, if f is reflection symmetric across N^\perp , then for every rotation P such that $PN = Y$, we must have that

$$\mathcal{F}_\ell = \mathcal{F}_\ell^* D_\ell(P)^* D_\ell(P)^\dagger. \quad (9)$$

This relationship helps to establish conditions for bilateral symmetry in the spherical harmonic domain. As discussed above, if $N = Y$, then $\mathcal{F}_\ell = \mathcal{F}_\ell^*$, and therefore the coefficients are real-valued, and we may write $\mathcal{F}_\ell = \mathcal{R}_\ell$ for real-valued vector \mathcal{R}_ℓ ; see Figure 3 for an example. If a function g is obtained by starting with a function f that is symmetric across Y^\perp , and rotating f by P , then we have that $\mathcal{G}_\ell = \mathcal{R}_\ell D_\ell(P)$. We call such functions *linear phase*. The terminology comes from the signal processing literature, because of the analogous situation on the real-line [5]. If a real-valued function h has symmetry across the origin, *i.e.*, $h(x) = h(-x)$ for all $x \in \mathbb{R}$, then its Fourier transform $H(\omega)$ is real-valued. Translation of such a function by x_0 results in $H(\omega) = r(\omega) e^{j\omega x_0}$, where $r(\omega)$ is a real-valued function for all frequencies ω . Such functions are described in signal processing as having linear phase, since the phase component ωx_0 is linearly dependent on the shift x_0 . In fact, if f is symmetric about a point $x_0 \in \mathbb{R}$, then it is easily shown that

$$F(\omega) = F(\omega)^* e^{-j\omega x_0} e^{-j\omega x_0} = F(\omega)^* e^{-2j\omega x_0}. \quad (10)$$

Note the similarity of (10) to (9), with the translation phase $e^{j\omega x_0}$ replaced by the rotation phase $D_\ell(P)$.

5. Finding the symmetry plane

Since symmetric functions have linear-phase spherical harmonic coefficients, the problem of finding the symmetry plane requires optimizing a linear phase fit to the observed coefficients. We formulate that optimization as follows.

Given observed coefficient vectors \mathcal{F}_ℓ for $\ell = 1, \dots, L$, we seek real-valued coefficient vectors $\{\mathcal{R}_\ell\}_{\ell=1}^L$ and a rotation P that jointly minimize the objective function

$$\Psi(\mathcal{R}, P) = \frac{\sum_{\ell=1}^L \|\mathcal{F}_\ell - \mathcal{R}_\ell D_\ell(P)\|^2}{\sum_{\ell=1}^L \|\mathcal{F}_\ell\|^2} \quad (11)$$

Since the generally suboptimal choice of setting all \mathcal{R}_ℓ to zero gives $\Psi = 1$, we conclude that, with the optimal choice, the minimum of Ψ must lie in $[0, 1]$, with lower values indicating a better linear-phase fit, and hence a more symmetric function.

Before optimizing (11), let us consider the degrees of freedom (DOF) between the data \mathcal{F}_ℓ and the model $\mathcal{R}_\ell D_\ell(P)$ that we are trying to fit, in order to ascertain when the fit is over, exactly, and under determined. Taking into account the conjugate symmetry $F_\ell^{-n} = (-1)^n F_\ell^{n*}$, we impose a similar symmetry for the real coefficients: $\mathcal{R}_\ell^{-n} = (-1)^n \mathcal{R}_\ell^n$. Enumerating the real degrees of freedom, we see that each \mathcal{F}_ℓ has $2\ell + 1$ real DOF, which we are trying to fit with $\ell + 1$ real parameters in \mathcal{R}_ℓ . In addition, we need to find one set of three Euler angles in P for all $\ell \leq L$. Hence, if we compare $\sum_{\ell=1}^L (2\ell + 1)$ real DOF in the data to $3 + \sum_{\ell=1}^L (\ell + 1)$ DOF in the model, then the fit is under-determined for $L = 1$, and exact for $L = 2$. For $L > 2$, we obtain a least-squares problem in which the data overdetermine the model.

For $L \geq 2$, the optimization of (11) proceeds as follows. By expanding the numerator for each ℓ , we obtain

$$\|\mathcal{F}_\ell\|^2 - \mathcal{R}_\ell D_\ell(P) \mathcal{F}_\ell^\dagger - \mathcal{F}_\ell D_\ell(P)^\dagger \mathcal{R}_\ell + \|\mathcal{R}_\ell\|^2. \quad (12)$$

Vector differentiation shows that the optimum \mathcal{R}_ℓ are

$$\mathcal{R}_\ell^{opt} = \frac{1}{2} (\mathcal{F}_\ell D_\ell(P)^\dagger + \mathcal{F}_\ell^* D_\ell(P)^\top). \quad (13)$$

Inserting \mathcal{R}_ℓ^{opt} back into (12) and simplifying, we obtain our main theoretical result: *the optimal choice for the rotation P that determines the orientation of the symmetry plane is obtained by maximizing the function*

$$\Phi(P) = \sum_{\ell=1}^L \text{Real} \{ \mathcal{F}_\ell D_\ell(P)^\dagger D_\ell^*(P) \mathcal{F}_\ell^\top \}. \quad (14)$$

Note that (14) is expressed only in terms of the spherical harmonics coefficients \mathcal{F} and the rotation P , and the previously unknown real coefficients \mathcal{R} have been eliminated.

Insight into the optimization of (14) is obtained by examining basic properties. First, due to the unitarity of the Wigner D matrices, we have an upper bound: $\Phi(P) \leq \sum_{\ell=1}^L \|\mathcal{F}_\ell\|^2$. Second, if f is symmetric across the x - z plane, then, as discussed above, \mathcal{F}_ℓ is real-valued. Consequently, we achieve the upper bound by setting $P = I$, as

expected. Third, if f is linear phase and therefore has coefficients $\mathcal{F}_\ell = \mathcal{R}_\ell D_\ell(P_N)$ where P_N rotates the x - z plane to N^\perp , then (14) becomes

$$\sum_{\ell=1}^L \text{Real} \{ \mathcal{R}_\ell D_\ell(P_N) D_\ell(P)^\dagger D_\ell^*(P) D_\ell^*(P_N)^\dagger \mathcal{R}_\ell^\top \}. \quad (15)$$

Setting $P = P_N$, the true symmetry plane, gives the global maximum in this case, as expected. Fourth, if P is a rotation about the y axis (Euler angles $\alpha = \gamma = 0$), then $D_\ell(P) = D_\ell(0, \beta, 0)$ is a real-valued matrix by (6). Consequently, for every rotation P_y about y axis, the function Φ is invariant: $\Phi(P_y P) = \Phi(P)$, as verified by inserting $P_y P$ into (14) and noting that $D_\ell(P_y)^\dagger D_\ell^*(P_y) = I$. Hence, optimizing Φ recovers only two DOF due to invariance. This is intuitively reasonable, since only two angles are required to specify the orientation of a plane in 3-D. Finally, calculation of (14) is efficient because it depends only on a relatively small number of spherical harmonic coefficients, and not on the much larger number of vertices. In our experiments we use a bandwidth of $L = 16$, resulting in only 288 coefficients, for shapes containing up to $N=160,000$ vertices.

6. Optimum estimates for the symmetry plane

The form of Φ in (14) may be rewritten using the trace identity $uAv^\top = \text{Trace}(v^\top uA)$ as follows

$$\Phi(P) = \sum_{\ell=1}^L \text{Real} \{ \text{Trace}(\mathcal{F}_\ell^\top \mathcal{F}_\ell D_\ell(P)^\dagger D_\ell^*(P)) \}. \quad (16)$$

In this form, Φ is tantalizingly close to the inverse Fourier transform on the rotation group $SO(3)$. The connection to $SO(3)$ is made more apparent if we define \tilde{P} to be the *conjugate rotation* such that $D_\ell(\tilde{P}) = D_\ell^*(P)$; if P has Euler angles (α, β, γ) , then \tilde{P} has Euler angles $(-\alpha, \beta, -\gamma)$. Then, setting $U = P^\top \tilde{P}$, we observe that $D_\ell(U) = D_\ell(P)^\dagger D_\ell^*(P)$. Therefore, if we define the objective function

$$\Omega(U) = \text{Real} \left\{ \sum_{\ell=1}^L \text{Trace}(\mathcal{F}_\ell^\top \mathcal{F}_\ell D_\ell(U)) \right\}, \quad (17)$$

we observe that both Ω and Φ reach the same maximum value. Optimization of Ω is straightforward by noting that the inner sum is exactly an inverse Fourier transform over the rotation group $SO(3)$, which suggests we apply the FFT developed for that group [8], and used for a similar purpose of correlation matching [9]. However, once we know the optimum value of U , we still need to obtain a P such that $P^\top \tilde{P} = U$. There are no simple equations for “inverting” U to obtain P , though an iterative numerical solution is possible. Furthermore, optimizing over U entails optimization




Figure 4: For the airplane model in (a), the three pseudocolor images in parts (b), (c), and (d) show respectively the cost function Φ as a function of α (horizontal) and β (vertical). Darker red values indicate improving estimates of the symmetry plane of (a). The global optimum, indicated with an arrow in (d), becomes clearer with increasing bandwidth.

over three Euler angles α , β , and γ , which is, in principle, unnecessary, as two angles suffice to specify the orientation of the symmetry plane.

We propose an alternative approach to optimization of Φ , by taking advantage of the property that $\Phi(P_y P_1) = \Phi(P_2)$ for each pair of rotations P_1, P_2 and every rotation P_y about the y -axis. The property suggests that optimization over only two of the three Euler angles, namely α , and β , are possible. Given a candidate rotation P , we compute $N = PY$, where $Y = [0, 1, 0]^\top$ as before. Given $N = [n_x, n_y, n_z]^\top$, the required angles are obtained as follows: $\beta = \text{atan2}(n_z, -n_x)$ and $\alpha = \text{acos}(n_y)$. Note that the Euler angle $\gamma = 0$. Given α, β , the cost of computing (6) for the Wigner matrix $D_\ell(\alpha, \beta, 0)$ is dominated by the cost of the computing the “little” Wigner d function $d_\ell(\beta)$. The d function is computed recursively for a given β , which is time-consuming. To reduce the computation, we take advantage of a decomposition [8], which shows that, as a function of the Euler angles:

$$P(\alpha, \beta, 0) = P(\alpha - \pi/2, -\pi/2, \beta) P(0, \pi/2, \pi/2). \quad (18)$$

This equation ensures that, once we compute $D_\ell(0, \pi/2, 0)$ offline for $\ell \leq L$, we may use those stored values to compute $D_\ell(P)$ efficiently from (6).

Figure 4 illustrates how the cost function Φ increases in detail with bandwidth L . For sufficiently high bandwidth, it is clear from the figure that the finding the optimum of the non-convex surface requires a search. To perform the search efficiently, we use a nested polar grid search, in which we

perform initially a coarse grid search by evaluating N values of α, β over $[0, \pi]$. (Note that α and $\alpha + \pi$ represent the same plane.) We then repeat the grid-search in a nested manner over the previously found best grid point, so that each repetition subdivides the previous best grid cell and its immediate neighbors as the domain of the new grid. The advantage of the nested search over, say, making the initial coarse grid much finer, is that we reject many grid cells in each repetition. We refer to this algorithm below as N-GRID. Note that FFT search methods [8][9] are essentially grid search methods, although they employ a fast algorithm to compute the grid values due to the objective being formulated as inverse $SO(3)$ transform. N-GRID is also similar to the recursive icosahedral search of Martinet *et al.* [11].

An alternative to grid search is stochastic search optimization. We use interacting simulating annealing [4] (ISA), a global optimization technique designed for multimodal objectives. We adapted the method to quaternion space by generating samples using the Fisher-Von Mises kernel:

$$f(x; \mu, \kappa) \propto \exp(\kappa \mu^T x), \quad (19)$$

where $\mu \in \mathbb{S}^3$ represent the mean and $\kappa \in \mathbb{R}$ is the concentration parameter. The distribution is transformed from a uniform distribution to a delta distribution as κ is changed from 0 to ∞ . We defined a fixed diffusion schedule such that, in first iteration $\kappa = 0$, and in subsequent iterations κ is defined by a power law. This ensures that the search is uniform over the entire manifold. We refer to this algorithm as SH-ISA below.

As a baseline for both algorithms, we use the three eigenvectors of the covariance matrix to obtain three corresponding candidates for the symmetry plane. The candidate that maximizes (14) is designed as the SH-COV estimate. We expect SH-COV to perform much worse than N-GRID and SH-ISA, due to its known failure mode when the covariance eigenvalues are nearly equal.

7. Experimental methods and results

We estimate spherical harmonic coefficients using the iterated residual fitting (IRF) method [3]. IRF calculates each vector \mathcal{F}_ℓ by using a weighted least squares approach based on the residual from lower frequencies, where the weight is adjusted using a smoothing parameter σ . We use $\sigma = 5 \times 10^{-4}$, which performs a mild amount of smoothing. We use a bandwidth of $L = 16$ for all shapes, as Figure 4 shows that it gives a clear peak at the global optimum. Evaluation of the objective function (14) takes less than 10ms on a computer with a 2.66 GHz quad-core processor and 10 Gb RAM.

Previous works have tested symmetry estimation on a relatively small number of shapes, and tested robustness by simulating acquisition and topological noise [1][11][20].




Figure 5: Uniform densification is illustrated for the plant model in (a), whose 316 vertices are plotted in (b), and whose densified samples are shown in (c).

While such simulations are worthwhile, there are now databases containing a wide variety of symmetric, nearly symmetric, and asymmetric shapes which provide a more extensive test. We use three databases: the Princeton Shape Benchmark (PSB) [18], the SHREC 2010 Generic 3-D warehouse (referred to as S10 here) [15] and the Mesh Segmentation Database (SEG) [2]. The PSB is a database of 1,814 shapes defined by triangulated meshes, while the S10 has 3,168, and SEG has 380. The shapes vary in the number of vertices from the hundreds to well over 100,000. However, the vertices are not uniformly distributed across the shape’s surface, which may affect the symmetry estimation. Therefore, we uniformly sample a dense set of points across the surface of the shape by relying on the triangulated mesh. For each triangle T , with vertices A, B , and C , a random point p lying inside the triangle may be generated using

$$p = (1 - \sqrt{r_1})A + \sqrt{r_1}(1 - r_2)B + \sqrt{r_1}r_2C, \quad (20)$$

where $r_1, r_2 \in [0, 1]$ are two independent, uniformly-distributed, random numbers [14]. The number of points in each triangle is calculated to be proportional to the triangle’s area in relation to the total area of all triangles. With such a representation, the points are more uniformly distributed on the shape surface, helping the identification of symmetry planes. This “densification” of vertices is illustrated in Figure 5.

In our experiments, every shape is randomly rotated in all three Euler angles prior to symmetry estimation. To provide a shape-independent measure of fit, we define the normalized fitting error as the relative error between the linear phase fit for the optimum rotation P_{opt} and its upper bound:

$$E(P_{opt}) = 100 \times \frac{\sum_{\ell=1}^L \|\mathcal{F}_\ell\|^2 - \Phi(P_{opt})}{\sum_{\ell=1}^L \|\mathcal{F}_\ell\|^2} \quad (21)$$

Low values of E , below 15%, indicate a good symmetry plane fit. Figure 6 shows examples from the PSB database of various degrees of symmetry and the associated fit. We found over numerous shapes that the E measure correlated well with our perception of how well the symmetry plane fit the shape.




Figure 6: The perfectly-symmetric airplane in (a) has nearly zero fitting error E . The middle legs of the nearly-symmetric insect in (b) are not matched, and the body is slightly twisted, raising error E . The branches of the asymmetric pine tree are not aligned, raising E to high levels. In each case the symmetry plane found by optimizing (14) with SH-ISA is shown superimposed in green.

Input \ Algo	SH-COV	N-GRID	SH-ISA
PSB EGI	29.1	16.2	15.7
PSB Vertex (O)	15.7	8.6	8.3
PSB Vertex (D)	8.8	5.6	5.3
S10 EGI	34.7	19.2	18.4
S10 Vertex (O)	19.7	10.4	9.9
S10 Vertex (D)	10.8	6.9	6.5
SEG EGI	35.7	19.5	18.7
SEG Vertex (O)	19.2	9.9	9.4
SEG Vertex (D)	9.6	6.1	5.7

Table 1: The normalized fitting error E is shown averaged over all shapes in each of three databases, using three algorithms. Rows 2-4 show results for the PSB database, and contain respectively the results for the EGI mapping, the vertex mapping using original vertices, denoted (O), and the vertex mapping using the densified vertices (D); similarly for other rows and S10 and SEG databases. For all three databases, densification significantly reduces error, and SH-ISA has the least error for all databases and mappings.

Table 1 summarizes the results on over 5,000 shapes in all databases. We find that the symmetry plane is significantly better estimated using the vertex mapping than the EGI mapping, reducing E by at least 30%, and also that densification improves the vertex mapping results substantially, by similar amounts. The EGI mapping is unique only for convex objects, and does not capture shape properties as well as vertex mapping (see Fig. 2). In our experiments, we used $N = 30$ for N-GRID, with up to 5 recursive evaluations, and use an average of 1,200 cost function evaluations for the SH-ISA algorithm. The N-GRID takes approximately 3 seconds per shape, and SH-ISA approximately 5 seconds, on the 2.66 GHz processor described above. The processing times are independent of the number of vertices




Figure 7: The symmetry plane from SH-ISA is shown superimposed in green. Asymmetries in the shapes, noticeable on magnification, result in nonzero E values.

in the shape, as a benefit of (14) relying on the spherical harmonic expansion alone. Densification allocates 100,000 samples to each surface, no matter how many vertices it has.

The robustness of our linear phase methods are due to the least-squares fit of the spherical harmonic coefficients using IRF, the ability of low-frequency spherical harmonics to smooth fluctuations, and performing optimization of (14) without differentiation. The large number of shapes tested here precludes showing many in the style of Figure 6, but Figure 7 shows representative results. More are presented as supplementary material for review purposes, along with the MATLAB code used.

The main limitation in symmetry determination is the spherical mapping. One important case is articulated shapes, which requires an analysis of intrinsic symmetry [17], and would normally fail with the extrinsic methods used here, as shown in the left column of Figure 8. However, by employing isometric embeddings which make the surface mesh roughly invariant to pose, we obtain much better planes of symmetry, as shown in the second column of Fig. 8. The third column shows the partition obtained from the embedding on the original surface.

8. Summary and conclusions

Figure 1 inspires us to look in the phase domain for information about symmetry. By developing the idea that linear phase in spherical harmonic coefficients is equivalent to bilateral symmetry, our main theoretical contribution shows that optimizing the degree of linear phase fit is equivalent to estimating the symmetry plane. The benefit of relying on spherical harmonic coefficients alone, as our method does,




Figure 8: The left column shows how symmetry detection fails on articulated shapes, the middle column shows the result after isometric embedding, and the right shows the result from the middle superimposed on the original shape.

is that they form a compact description of a shape; indeed, that compactness is a primary motivator to their many uses in 3-D modeling and vision [3][7][9][10]. Our method is compatible with those works in that it is able to use any symmetry-preserving mapping to the sphere, including the vertex mapping and EGI mapping.

Beyond symmetry, our results suggest that the phase of spherical harmonics is rich in other information about 3-D structure. It is possible that non-rigid symmetry, such as spiral structure, is connected to polynomial phase, rather than linear phase. Furthermore, the structural properties of phase for three-variable spherical harmonics, or SPHARM [16] constitute another interesting extension to explore.

References

- [1] A. Bermanis, A. Averbuch, and Y. Keller. 3-D symmetry detection and analysis using the pseudo-polar Fourier transform. *International Journal of Computer Vision*, 90(2):166–182, 2010. 2, 6
- [2] X. Chen, A. Golovinskiy, and T. Funkhouser. A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 28(3), Aug. 2009. 6
- [3] M. K. Chung, K. M. Dalton, L. Shen, A. C. Evans, and R. J. Davidson. Weighted Fourier series representation and its application to quantifying the amount of gray matter. *IEEE Transactions on Medical Imaging*, 26(4):566–581, 2007. 1, 3, 6, 8
- [4] J. Gall, J. Potthoff, C. Schnrr, B. Rosenhahn, and H.-P. Seidel. Interacting and annealing particle filters: Mathematics and a recipe for applications. *Journal of Mathematical Imaging and Vision*, pages 1–18, 2007. 6
- [5] R. Kakarala and J. A. Cadzow. Estimation of phase for noisy linear phase signals. *IEEE Transactions on Signal Processing*, 44(10):2483–2497, 1996. 4
- [6] M. M. Kazhdan, B. Chazelle, D. P. Dobkin, A. Finkelstein, and T. A. Funkhouser. A reflective symmetry descriptor. In *ECCV (2)*, pages 642–656, 2002. 1, 2
- [7] M. M. Kazhdan, T. A. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Symposium on Geometry Processing*, pages 156–164, 2003. 1, 8
- [8] P. Kostelec and D. Rockmore. FFTs on the rotation group. *Journal of Fourier Analysis and Applications*, 14:145–179, 2008. 5, 6
- [9] A. Makadia and K. Daniilidis. Spherical correlation of visual representations for 3D model retrieval. *International Journal of Computer Vision*, 89(2-3):193–210, 2010. 1, 2, 3, 5, 6, 8
- [10] A. Makadia, A. Patterson, and K. Daniilidis. Fully automatic registration of 3D point clouds. In *CVPR*, pages 1297–1304, 2006. 1, 8
- [11] A. Martinet, C. Soler, N. Holzschuch, and F. Sillion. Accurate detection of symmetries in 3D shapes. *ACM Transactions on Graphics*, 25(2):439 – 464, April 2006. 1, 2, 3, 6
- [12] N. J. Mitra, L. J. Guibas, and M. Pauly. Partial and approximate symmetry detection for 3D geometry. *ACM Transactions on Graphics*, 25(3):560–568, 2006. 1
- [13] N. J. Mitra, M. Pauly, M. Wand, and D. Ceylan. Symmetry in 3D geometry: Extraction and applications. In *EUROGRAPHICS State-of-the-art Report*, 2012. 2
- [14] R. Osada, T. A. Funkhouser, B. Chazelle, and D. P. Dobkin. Shape distributions. *ACM Transactions on Graphics*, 21(4):807–832, 2002. 6
- [15] V. T. Porethi, A. Godil, H. Dutagaci, T. Furuya, Z. Lian, and R. Ohbuchi. SHREC’10 track: Generic 3D warehouse. In *3DOR*, pages 93–100, 2010. 6
- [16] M. Quicke, C. Brechbühler, J. Hug, H. Blattman, and G. Székely. Parameterization of closed surfaces for parametric surface description. In *CVPR*, pages 1354–1360, 2000. 8
- [17] D. Raviv, A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Full and partial symmetries of non-rigid shapes. *International Journal of Computer Vision*, 89(1):18–39, 2010. 2, 7
- [18] P. Shilane, P. Min, M. M. Kazhdan, and T. A. Funkhouser. The princeton shape benchmark. In *SMI*, pages 167–178, 2004. 6
- [19] S. N. Sinha, K. Ramnath, and R. Szeliski. Detecting and reconstructing 3D mirror symmetric objects. In *ECCV (2)*, pages 586–600, 2012. 2
- [20] C. Sun and J. Sherrah. 3D symmetry detection using the extended gaussian image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):164–168, 1997. 1, 2, 6
- [21] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii. *Quantum theory of angular momentum*. World Scientific, Singapore, 1988. 2