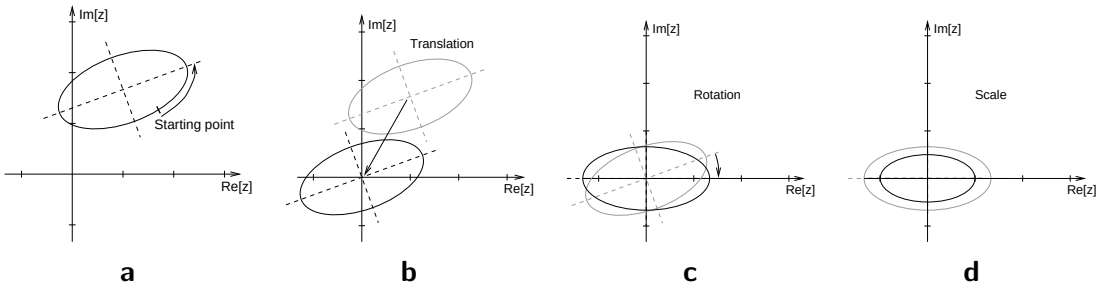

Figure 4.2: Normalization steps of Fourier coefficients; shifting of the starting point to the tip of the ellipse (**a**), moving the center of gravity to the coordinate origin (**b**), rotating the main axis of the ellipse to the real axis (**c**), and finally scaling the half major axis to unity (**d**).

where the real valued coefficients a_n, b_n, c_n, and d_n are defined as follows.

$$
\begin{aligned}
a_n &= \mathrm{Re}z_n + \mathrm{Re}z_{-n} && (4.17)\\
b_n &= -\mathrm{Im}z_n + \mathrm{Im}z_{-n} && (4.18)\\
c_n &= \mathrm{Im}z_n + \mathrm{Im}z_{-n} && (4.19)\\
d_n &= \mathrm{Re}z_n - \mathrm{Re}z_{-n} && (4.20)\\
& && (4.21)
\end{aligned}
$$

4.2 Description of surfaces by spherical harmonic functions

The problem of finding a similarly homogeneous parametrization of arbitrarily shaped surfaces proved to be more difficult. [Brechbühler *et al.* 1995] introduced only recently a new surface parametrization technique which can be considered as a complete generalization of Kuhl and Giardina's technique to three dimensions.

Earlier methods for mapping an object surface onto a sphere have been limited to represent only star-shaped or convex objects, as they start from an initial radial surface function $r(\theta, \phi)$. [Staib and Duncan 1992b] discuss the use of a parameter space with torus topology, which can be deformed into a tube by squeezing the torus cross-section to a thin ribbon. Closed surfaces are obtained by considering tubes whose ends close up to a point. This approach illustrates some principal difficulties which can also be found in other parametrization techniques.

- Warping a torus to a closed surface poses the problem that the *parameters have different rules*. One parameter defines a kind of spine along which cross-sections are stacked up.
- Squeezing a circle to line results in a *nonhomogeneous distribution of parameters* on the object surface.
- Warping a torus to a tube and finally to a closed surface causes the parametrization does *not* result in a *one-to-one mapping* of surface points to parameters.

The new method allows a uniform mapping of an object surface into a two-coordinate space with spherical topology. As a mapping of convoluted surface structures onto the surface of a sphere introduces distortions, optimization of the distribution of nodes in parameter space becomes necessary. This problem is solved by nonlinear optimization.

Parametrized surfaces can be expanded into spherical harmonics, hierarchically describing shape properties by spatial frequency constituents. Similarly to Fourier descriptors, spherical harmonic descriptors can also be made invariant to translation, rotation and scaling.

A summary of the surface parametrization procedure follows which is described in detail in [Brechbühler *et al.* 1992, Brechbühler *et al.* 1995].

4.2.1 The surface data structure

Medical CT or MRI images are examples of volumetric data. For each cuboidal cell (volume element or voxel) in a certain volume there are one or more measurements. When segmentation succeeds, one anatomical unit can be characterized by a binary data volume, in which every voxel contains either 1, which means it belongs to the unit, or 0, meaning it is in the background. The object is then the set of "1" voxels and can be pictured as a collection of small cubes, adopting the *cuberille notion* [Herman and Liu 1979]. The surface of a voxel object is a set of unit squares, all parallel to one of the three coordinate planes yz, xz, or xy. The edges and vertices that bound the faces are also parts of the surface, which is represented as a data structure that reflects geometry as well as neighborhood relations.

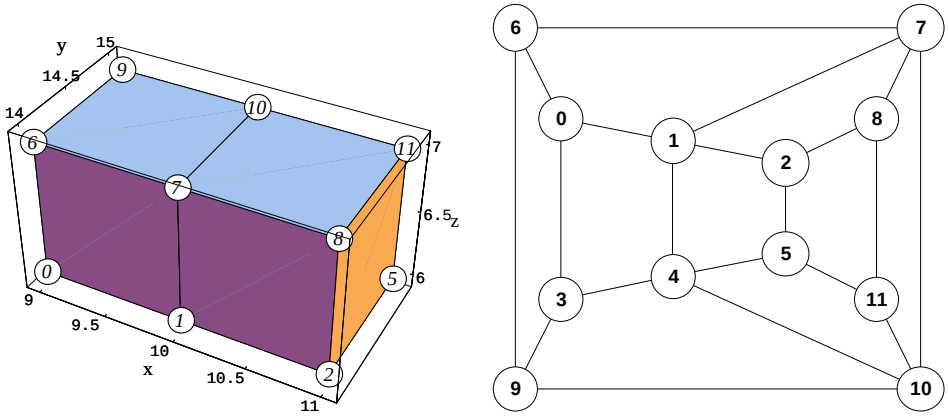

Figure 4.3: Left: A two-voxel example object illustrating the surface data structure, which focuses on the vertices. The numbers of the vertices are shown in circles (vertices 3 and 4 are hidden). The data structure entry for a vertex represents its Cartesian coordinates and a list of neighboring nodes. The entry for vertex 7 is {{x_0 =10, x_1 =14, x_2 =7}, *neighbors*={6, 0, 1, 2, 8, 11, 10, 9}}. The table below lists the complete surface data structure of this object. — **Right:** A flat diagram of the surface net for the same object. (courtesy of Ch. Brechbühler)

node nr.	x_0	x_1	x_2	neighbors
0	{{{	9,	14, 6},	{ 1, 7, 6, 9, 3, 4}},
1	{{	10,	14, 6},	{ 0, 3, 4, 5, 2, 8, 7, 6}},
2	{{	11,	14, 6},	{ 1, 4, 5, 11, 8, 7}},
3	{{	9,	15, 6},	{ 4, 1, 0, 6, 9, 10}},
4	{{	10,	15, 6},	{ 3, 9, 10, 11, 5, 2, 1, 0}},
5	{{	11,	15, 6},	{ 4, 10, 11, 8, 2, 1}},
6	{{	9,	14, 7},	{ 7, 10, 9, 3, 0, 1}},
7	{{	10,	14, 7},	{ 6, 0, 1, 2, 8, 11, 10, 9}},
8	{{	11,	14, 7},	{ 7, 1, 2, 5, 11, 10}},
9	{{	9,	15, 7},	{ 10, 4, 3, 0, 6, 7}},
10	{{	10,	15, 7},	{ 9, 6, 7, 8, 11, 5, 4, 3}},
11	{{	11,	15, 7},	{ 10, 7, 8, 2, 5, 4}}}

Table 4.2: The complete surface data structure of the two cube object.

4.2.2 Parametrization of closed surfaces

A key step in the shape description of a surface is its mapping to the parameter space, the sphere Ω_3. Any point on the surface must map to exactly one point on the sphere, and vice versa. The location on the sphere corresponding to a surface point defines the *parameters* of the point. It can be represented as two polar or three Cartesian coordinates, related through the bijection

$$\begin{pmatrix} u_0 \\ u_1 \\ u_2 \end{pmatrix} = \begin{pmatrix} \sin\theta\,\cos\phi \\ \sin\theta\,\sin\phi \\ \cos\theta \end{pmatrix} .$$

Mapping a surface to the sphere assigns parameters to every surface point; therefore it is called *surface parametrization*. The mapping must be continuous, i.e. neighboring points in one space must map to neighbors in the other space. It is possible and desirable to construct a mapping that preserves areas (see Figure 4.8). Narrowing to the cuberille notion, Figure 4.4 symbolically illustrates this mapping of a selected facet from the object surface to a portion of Ω_3. It is not possible in general to map every surface facet to a spherical square: distortions cannot be avoided, but they should be minimal.

The parametrization, i.e., the embedding of the object surface graph into the surface of the unit sphere can be formulated as a constrained optimization problem. The following paragraphs define the meaning of *variables*, *objective* (goal function), *constraints* and *starting values* in this context.

Variables

The coordinates of all vertices vary in the optimization. Using two (e.g. spherical) coordinates per vertex would be the most economic representation with respect to storage space, but it would make the equal treatment of all spatial directions difficult and pose

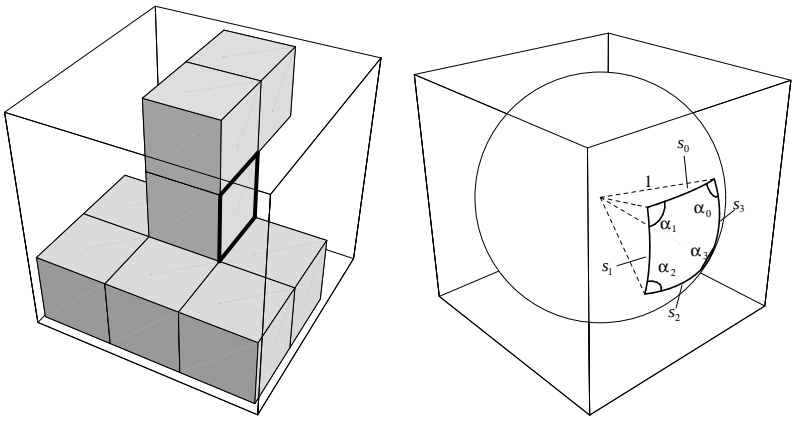

Figure 4.4: Every single face on the object's surface is mapped to a spherical quadrilateral. The sides of a spherical polygon are geodesic arcs on the sphere surface. As the sphere has unity radius, the length of a side s_i is equal to the corresponding center angle (in radian). The quadrilateral in this illustration is special in that its four sides $s_0 \cdots s_3$ are equal and its four angles $\alpha_0 \cdots \alpha_3$ are equal: it is the spherical analogue of a square. (courtesy of Ch. Brechbühler)

the problem of discontinuity and singularities in the parameter space. The authors prefer Cartesian coordinates (u, v, w) for representing a location on the sphere, introducing one virtual degree of freedom per vertex. The number of variables is three times the number of vertices.

Constraints

Two kinds of equalities and one kind of inequality constrain the values that the variables can take.

1. Every vertex must lie on the unit sphere in parameter space, i.e., $u_0^2 + u_1^2 + u_2^2 = 1$. This constraint compensates for the virtual degree of freedom and forces.

2. Ask for *area preservation*; any object surface region must map to a region of proportional area on the sphere. There is one constraint included for each elementary facet: the area of the spherical quadrilateral must be 4π divided by the total number faces.

3. All quadrilaterals on the sphere must remain convex; no angle α_k may become negative or exceed π.

Objective

The objective is to minimize the distortion of the surface net in the mapping. It is conceptually similar to angle preservation, and it must tend to make the shape of all the mapped faces as similar to their original square form as possible. To fulfill this goal perfectly, a facet should map to a "spherical square" (see Figure 4.4). This can never be reached exactly for all faces and a trade off has to be made between the distortions at different vertices. To achieve an optimum, the goal function maximizes $\sum_{i=0}^{3} \cos s_i$

summed over all spherical quadrilateral, which is the same as the sum of the cosines of the lengths of all edges.

Starting values

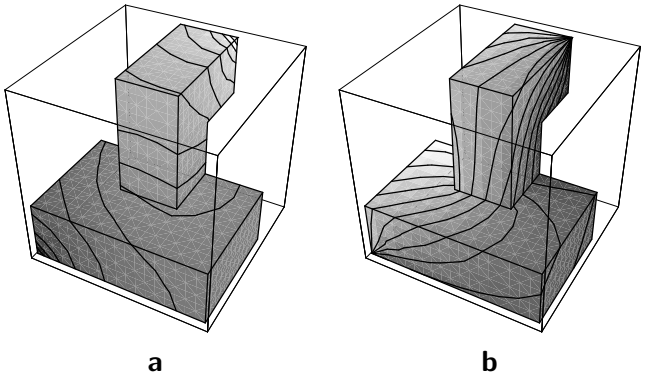

Figure 4.5: The simple object "duck" consisting of nine voxels is used for illustrating the initial parametrization. The north pole is at the lower left, the south pole at the upper right. Co-latitude is mapped on the object's surface as a grey value in *a*; iso-latitude lines are drawn every $\frac{\pi}{16}$. Longitude is shown in *b*; iso-longitude lines ("meridians") are $\frac{\pi}{8}$ apart. (courtesy of Ch. Brechbühler)

The variables of the optimization are the positions on the unit sphere to which the vertices are mapped. Therefore, starting values in this context means an initial mapping of the object's surface to the sphere. Two vertices are selected as the poles of the spherical parameter space and a two step diffusion process is applied to assign latitude and longitude coordinates to each vertex (see Figure 4.5).

Optimization method

The commonly available optimization routines can not be used for the resulting nonlinear constrained minimization problem of such a large scale, as they do not exploit its sparsity and information available about the constraints. A Newton-Lagrange algorithm has been developed by the author to find the constrained minimum of the goal function.

The solution of the nonlinear program defines the optimal parametrization of the object's surface. Figure 4.6 shows the starting point for the optimization; the right diagram is a combination of those in Figure 4.5. Figure 4.7 visualizes the result of the optimization in different ways. The same vertex as in Figure 4.6 is marked. Figure 4.8 demonstrates the importance of homogeneous surface parametrization.

4.2.3 Parametrization by spherical harmonics basis functions

The description of the surfaces of simply connected 3D objects in an arbitrary basis can be performed similarly to the 2D case. The surface will be parametrized by two variables, the θ and ϕ polar parameters, and will be defined by three explicit functions

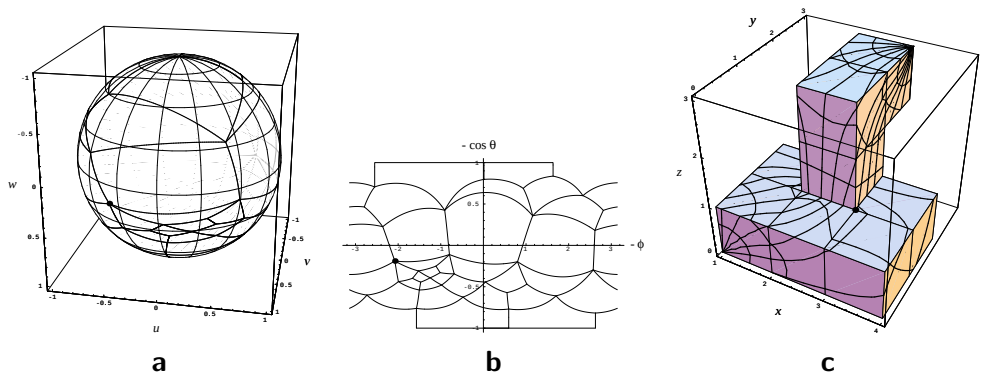

Figure 4.6: Diffusion yields the initial parametrization, which is plotted in the same three ways as the final result in Figure 4.7. **a** The surface net is plotted on the spherical parameter space. The thick lines depict the edges of the original square faces. The equidistance for both θ and ϕ is $\frac{\pi}{8}$. **b** ϕ and $\cos\theta$ are interpreted as Cartesian coordinates. The monotonic cosine function is applied to give a true-area cylindrical projection. The horizontal lines at ± 1 are the poles. **c** Conversely, the globe coordinate grid is drawn over the object. (courtesy of Ch. Brechbühler)

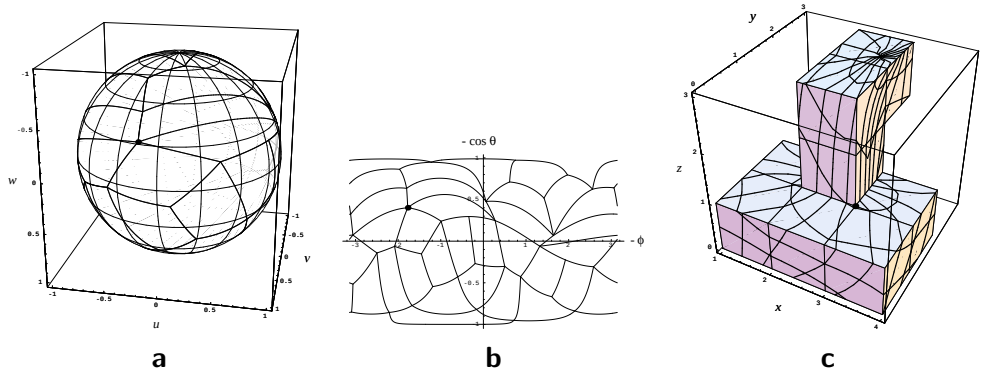

Figure 4.7: The result of the optimization, plotted in the same ways as in the previous figure. The areas of elementary facets in parameter space are now equal, and local distortions are minimized. The rotational position of the net on the sphere is arbitrary. **c** After optimization, the former poles have lost their prominent role. They have now the same importance as other point in parameter space and could lie anywhere on the surface. Only the use of polar coordinates for visualization gives them a conspicuous appearance. (courtesy of Ch. Brechbühler)

$$\boldsymbol{v}(\theta,\phi) = \begin{pmatrix} x(\theta,\phi) \\ y(\theta,\phi) \\ z(\theta,\phi) \end{pmatrix} \quad . \tag{4.22}$$

It has to be emphasized that this is not a radial function. Selecting the spherical harmonic functions (Y_l^m denotes the function of degree l and order m, see [Greiner and Diehl 1986]) as a basis, the coordinate functions can be written as

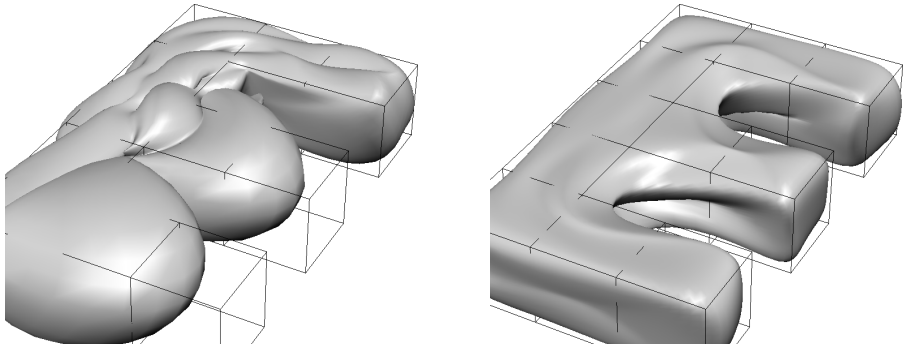

Figure 4.8: Homogeneous parameter distribution is important for shape description. The "E"-shaped object surface, indicated by a wireframe, is expanded into a series of spherical harmonics. The initial, non-uniform parametrization yields a poor shape representation (**a**); its optimization achieves a significant improvement (**b**). (courtesy of Ch. Brechbühler)

$$\boldsymbol{v}(\theta,\phi,\boldsymbol{p}) = \sum_{k=0}^{K} \sum_{m=-k}^{k} \boldsymbol{c}_k^m \; Y_k^m(\theta,\phi) \qquad , \tag{4.23}$$

where

$$\boldsymbol{c}_k^m = \begin{pmatrix} c_{x_k}^m \\ c_{y_k}^m \\ c_{z_k}^m \end{pmatrix} \qquad . \tag{4.24}$$

The expansion is restricted to the first $K+1$ terms. When the free variables θ and ϕ run over the whole sphere (e.g. $\theta = 0 \ldots \pi$, $\phi = 0 \ldots 2\pi$), the point $\boldsymbol{v}(\theta,\phi)$ runs over the whole surface of the object. The sphere Ω_3 is considered a perfectly symmetric surface without any singular points or preferred directions.

The surface is then described by the parameters

$$\begin{aligned} \boldsymbol{p} \;\; = \;\; & (c_{x0}^{\;0}, c_{y0}^{\;0}, c_{z0}^{\;0}, c_{x1}^{\;-1}, c_{x1}^{\;0}, c_{x1}^{\;1}, c_{y1}^{\;-1}, c_{y1}^{\;0}, c_{y1}^{\;1}, \\ & c_{z1}^{\;-1}, c_{z1}^{\;0}, c_{z1}^{\;1}, \ldots c_{xK}^{\;-K} \ldots c_{zK}^{\;K})^\top \end{aligned}$$

Y_l^m denotes the spherical harmonic function of degree l and order m. The following definitions agree with [Press *et al.* 1988], page 252. The variable w is a scalar and will correspond to u_2 below.

Legendre polynomials

$$P_l(w) \;\; = \;\; \frac{1}{2^l l!} \; \frac{d^l}{dw^l} \, (w^2-1)^l \tag{4.25}$$

Associated Legendre polynomials

$$\begin{aligned} P_l^m(w) \;\; &= \;\; (-1)^m (1-w^2)^{\frac{m}{2}} \frac{d^m}{dw^m} P_l(w) \\ &= \;\; \frac{(-1)^m}{2^l l!} (1-w^2)^{\frac{m}{2}} \frac{d^{m+l}}{dw^{m+l}} (w^2-1)^l \end{aligned} \tag{4.26}$$

Y_l^m	$m=0$	$m=1$	$m=2$
$l=0$	$\frac{1}{2\sqrt{\pi}}$		
$l=1$	$\sqrt{\frac{3}{4\,\pi}}\,\cos(\theta)$	$-\sqrt{\frac{3}{8\,\pi}}\,e^{i\,\phi}\,\sin(\theta)$	
$l=2$	$\sqrt{\frac{5}{16\,\pi}}\,\left(-1+3\cos(\theta)^2\right)$	$-\sqrt{\frac{15}{8\,\pi}}\,e^{i\,\phi}\,\cos(\theta)\,\sin(\theta)$	$\sqrt{\frac{15}{2\,\pi}}\,e^{2\,i\,\phi}\sin(\theta)^2$

Table 4.3: Spherical harmonic basis functions up to degree 2.

Spherical harmonic functions

$$Y_l^m(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi}\,\frac{(l-m)!}{(l+m)!}}\;\; P_l^m(\cos\theta)\;\; e^{im\phi} \tag{4.27}$$

$$Y_l^{-m}(\theta,\phi) = (-1)^m Y_l^{m*}(\theta,\phi) \tag{4.28}$$

A list of the spherical harmonics up to degree 2 in table 4.3 exemplifies these definitions in polar coordinates (θ,ϕ). For convinience, the real and imaginary parts of the complex basis functions Y_l^m can be used as independent real valued basis functions. There is the same number of functions as in the complex basis, namely $2l+1$ for any non-negative l.

Figures 4.9 and 4.9 provide two different ways to visualize the basis functions. In figure 4.9 function values of $\mathrm{Re}\,Y_l^m$ up to degree 3 are mapped onto the surface of a unit sphere representing the parameter space in θ and ϕ, light areas corresponding to positive values and dark to negative values. In physics spherical harmonics are often shown as in figure4.9 where the radius of a unit sphere is modulated by the function values $|\mathrm{Re}\,Y_l^m|$.

One has to realize that the parametrization by spherical harmonics is just one possibility for the parametric description of contours. Alternative methods, as e.g. deformable superquadrics have also been proposed in the literature (see previous chapter). The author emphasizes that any reasonable parametric shape model can be used within the presented formalism.

4.2.4 Invariant descriptors

The coefficients obtained so far still depend on the relative position of the parameter net of the object surface, on the orientation of the object space, on the size of the object and finally on its position in space. By transforming the object to canonical positions in parameter space and object space, similarly as it has been done with Fourier parametrized 2-D contours, one can get rid of these dependencies.

Rotation independent descriptors

Rotation of the object to standard position in parameter space and object space needs three rotations in each spaces, when rotations are described by Euler angles. We wish to rotate the object in parameter space so that the north pole ($\theta=0$, on the u_2 axis) will be at one end of the shortest main axis of this first order ellipsoid and the point where the Greenwich meridian ($\phi=0$) crosses the equator ($\theta=\frac{\pi}{2}$, on the u_0 axis) is at one end of

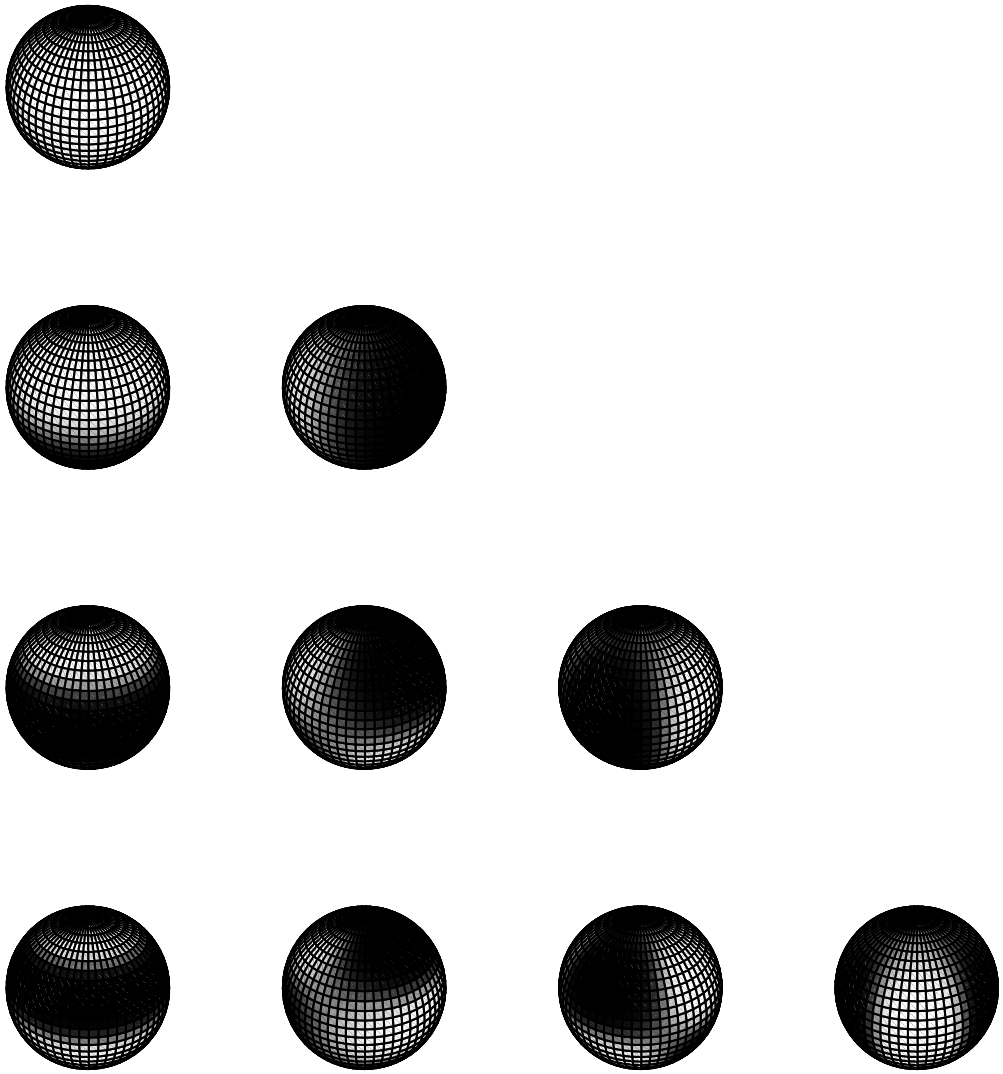

Figure 4.9: Real parts of spherical harmonic functions Y_l^m up to degree 3 are mapped onto the surface of a unit sphere, light areas corresponding to positive values and dark to negative values.

the longest main axis. To define a standard position only the contribution of the spherical harmonics of degree $l = 1$ in equation 4.23 are taken into consideration, these define an ellipsoid.

$$\boldsymbol{v}_1(\theta, \phi, \boldsymbol{p}) = \sum_{m=-1}^{1} \boldsymbol{c}_1^m \, Y_1^m(\theta, \phi) \tag{4.29}$$

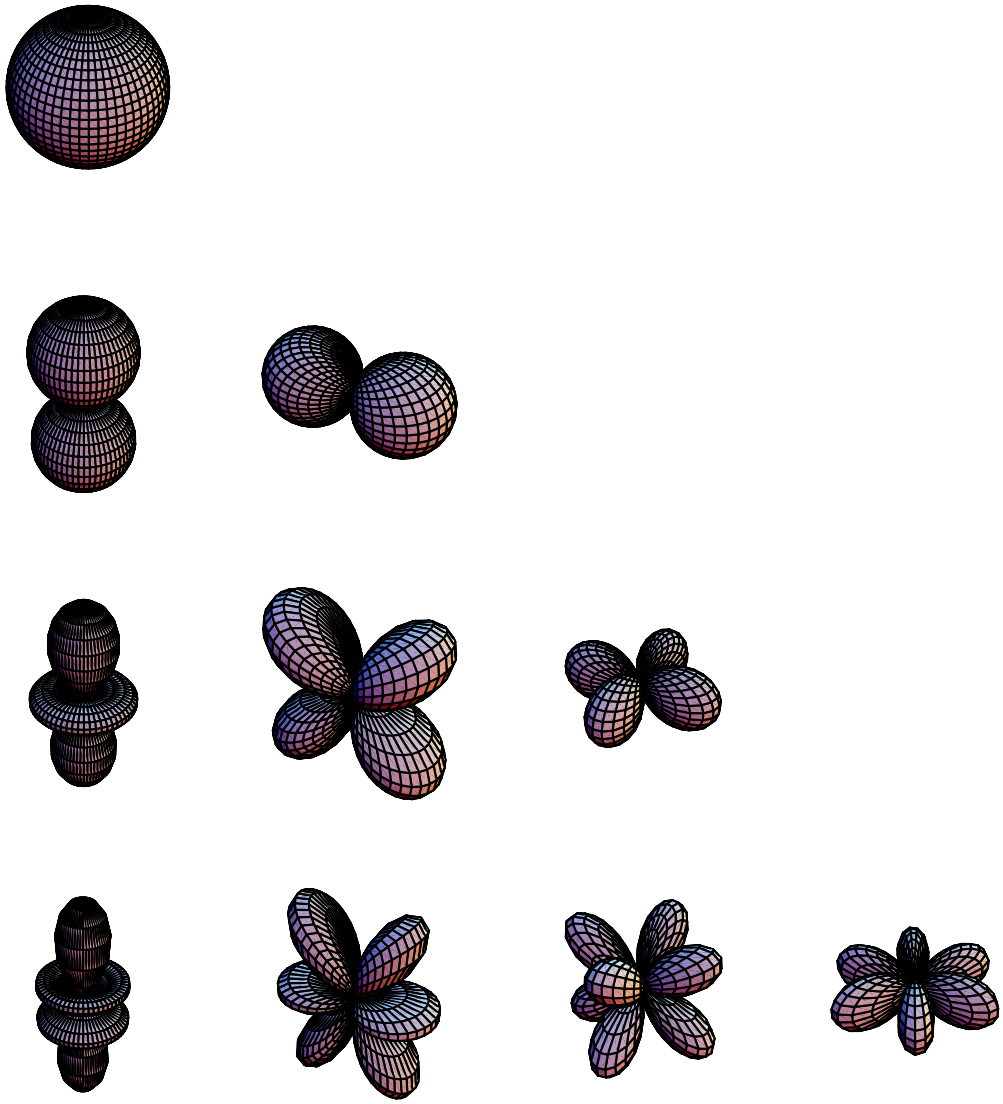

Figure 4.10: Real parts of spherical harmonic functions Y_l^m up to degree 3. The radius of a unit sphere is modulated by the function values $|\mathrm{Re}\, Y_l^m|$

Substituting the basis functions $Y_1^{-1} = \frac{\sqrt{3}}{2\sqrt{2\pi}}\,(u_0 - iu_1)$, $Y_1^0 = \frac{\sqrt{3}}{2\sqrt{\pi}}u_2$ and $Y_1^1 = -\frac{\sqrt{3}}{2\sqrt{2\pi}}(u_0 + iu_1)$ this sum can be written as

$$\boldsymbol{v}_1(\boldsymbol{u}) \;=\; \boldsymbol{A}\boldsymbol{u} \;=\; \boldsymbol{A}\begin{pmatrix} u_0 \\ u_1 \\ u_2 \end{pmatrix} \;=\; \boldsymbol{a}_1\, u_0 + \boldsymbol{a}_2\, u_1 + \boldsymbol{a}_3\, u_2 \quad , \tag{4.30}$$

where

$$\boldsymbol{A} = (\boldsymbol{a}_1, \boldsymbol{a}_2, \boldsymbol{a}_3) \quad = \quad \frac{\sqrt{3}}{2\sqrt{2\pi}} \left(\boldsymbol{c}_1^{-1} - \boldsymbol{c}_1^1, i(\boldsymbol{c}_1^{-1} + \boldsymbol{c}_1^1), \sqrt{2}\boldsymbol{c}_1^0 \right) \quad . \tag{4.31}$$

The rotation matrix $\boldsymbol{R}_u^T = (\hat{\boldsymbol{u}}_1, \hat{\boldsymbol{u}}_2, \hat{\boldsymbol{u}}_3)$, where $\hat{\boldsymbol{u}}_1$, $\hat{\boldsymbol{u}}_2$, and $\hat{\boldsymbol{u}}_3$ denote the unit eigenvectors of $\boldsymbol{A}$, can be then used to rotate the parameter net over the surface into the desired location. The roots of the eigenvalues $l_1^2 > l_2^2 > l_3^2$ represent half the length of the main axes of the ellipsoid. The rotation matrix is applied to the parameters $\boldsymbol{u}_i$ associated with each vertex i:

$$\boldsymbol{u}_i \mid^V = \boldsymbol{R}_u^T veu_i \ . \tag{4.32}$$

This new parametrization results in new coefficients $\boldsymbol{c}_l^m \mid^V$.

In the next step, the ellipsoid is rotated in the object space to make its main axes coincide with the coordinate axes, putting the longest ellipsoid axis along x and the shortest along z. This requires only the matrix multiplication

$$\boldsymbol{c}_l^m \mid^R = \boldsymbol{R}_x \ \boldsymbol{c}_l^m \mid^V \ , \tag{4.33}$$

where the rotation matrix is defined as $\boldsymbol{R}_x = \text{diagonal}(\frac{1}{l_1}, \frac{1}{l_2}, \frac{1}{l_2})\boldsymbol{A}^T \boldsymbol{R}_u^T$.

Scale independence

Scaling invariance can be achieved by dividing all descriptors by l_1, the length of the longest main axis

$$\boldsymbol{c}_l^m \mid^S = \boldsymbol{R}_x \ \boldsymbol{c}_l^m \mid^R \ , \tag{4.34}$$

Invariant spherical harmonic descriptors

Ignoring the coefficients of degree $l = 0$, that is setting $\boldsymbol{c}_0^0 \mid^T = (0, 0, 0)^T$ achieves translation invariance.

4.3 Ambiguous cases of normalization

The normalization techniques described here for the 2-D and 3-D cases require the precondition that coefficients of degree 1 represent a real ellipse or a real ellipsoid, respectively. If, however, the ellipse degenerates to a circle and the ellipsoid to an ellipsoid of revolution or a sphere, the technique will fail to derive stabile main axes. Objects of higher symmetries, such as regular polygons and polyhedra are a good example for the limitations of the normalization technique (see Figure 4.11). In general, the method will fall through for shapes with properties which are not reflected in the first degree cofficients. In these cases other features have to be incorporated into the normalization. These features can involve, e.g., coefficients of higher degree or an external coordinate system. Wyskočil in his diploma work [Wyskocil 199798] proposes a method which uses higher degree Fourier coefficients to unambiguously normalize 2-D contours.

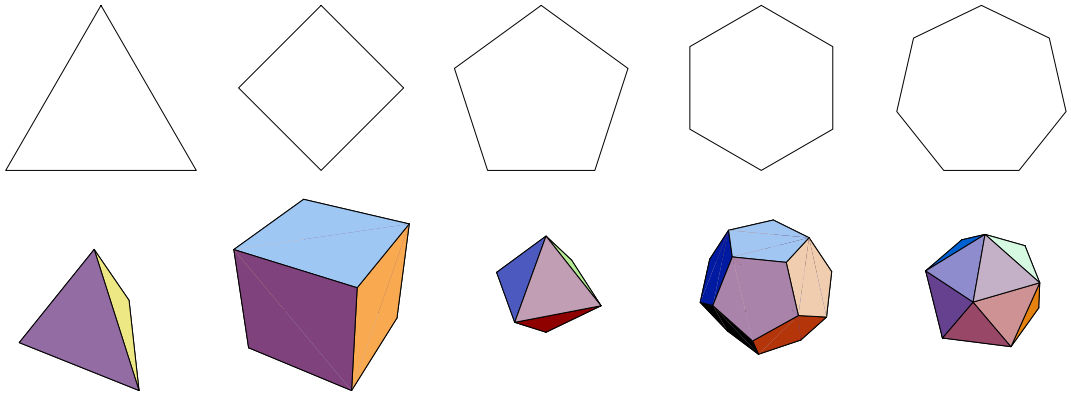

Figure 4.11: Because of the limitations of the normalization technique based solely on the first degree coefficients objects of higher symmetries, e.g. regular polygons and polyherda cannot be unambiguously aligned, since their fitst ellipse or ellipsoid degenerates to a circle or sphere.

The symmetries of the ellipse and ellipsoid have also to be taken into consideration while computing corresponding descriptors of similar objects. We will refer to this problem in the next chapters when we describe the generation of average models of anatomical organs.