

Faithful Least-Squares Fitting of Spheres, Cylinders, Cones and Tori for Reliable Segmentation

Gabor Lukács^{1,2}, Ralph Martin², and Dave Marshall²

¹ Computer and Automation Research Institute, Hungarian Academy of Sciences,
H-1518 Budapest, POB 63, Hungary

lukacs@sztaki.hu

² Dept of Computer Science, Cardiff University, PO Box 916, Cardiff, UK, CF2 3XF
ralph@cs.cf.ac.uk, dave@cs.cf.ac.uk

Abstract. This paper addresses a problem arising in the reverse engineering of solid models from depth-maps. We wish to identify and fit surfaces of known type wherever these are a good fit. This paper presents a set of methods for the least-squares fitting of spheres, cylinders, cones and tori to three-dimensional point data. Least-squares fitting of surfaces other than planes, even of simple geometric type, has been little studied.

Our method has the particular advantage of being robust in the sense that as the principal curvatures of the surfaces being fitted decrease (or become more equal), the results which are returned naturally become closer and closer to those surfaces of 'simpler type', *i.e.* planes, cylinders, cones, or spheres which best describe the data, unlike other methods which may diverge as various parameters or their combination become infinite.

1 Introduction

The motivation for this paper lies in reverse engineering of the geometric shape of simple mechanical components. A laser scanner is used to capture three-dimensional point data from the surface of an object. From this we wish to construct a boundary representation solid model of the object's shape. In particular, we wish to identify and fit particular simple surfaces to portions of the boundary wherever these are in good agreement with the point data. The problem can be decomposed into two logical steps: *segmentation*, where the data points are grouped into sets each belonging to a different surface, and *fitting*, where the best surface of an appropriate type is fitted to each set of points. The new results in this paper concern the latter problem, but we first outline the segmentation method used, as it has a significant effect on the final model created. We present new results on the fitting of spheres, cylinders, cones and tori, after giving a review of previous approaches. Detailed analysis and explicit formulae for the partial derivatives used in the non-linear least-squares optimisation can be found in [16].




Fig. 1. Initial 3D data; seed placement

2 Segmentation

Segmentation is the problem of grouping the points in the original dataset into subsets each of which logically belongs to a single primitive surface. Various approaches exist for segmenting simple surfaces from three-dimensional data [1–3, 5, 7, 10, 15]. Most commonly, segmentation is treated as a local-to-global aggregation problem with similarity constraints employed to control the process. Often several stages are required, mostly applied in a sequential fashion, ranging from the estimation of local surface properties such as curvature, to more complex feature clustering such as symmetry seeking. Typically, small initial seed regions are chosen at random positions. These are then grown and homogeneous regions are merged together. However, such approaches tend to isolate the segmentation stage from the representation stage with the result that the data partitioning may not agree well with the given primitive types. In addition, the sensitivity of these methods to noise in the data (especially outliers) may also lead to misclassification and, hence, poor results [12]. An efficient and reliable segmentation process thus depends on employing geometric knowledge of the primitive types, firstly, to guide the detection and grouping processes, and secondly, to assure the coherence and consistency of models throughout the whole segmentation process [1].

The approach used for segmentation in our reverse engineering project is the *recover and select* paradigm of Leonardis et al. [11, 13, 12, 14]. Here an iterative process recovers and selects specific instances of the required geometric primitives (planes, spheres, cylinders, cones and tori). The basic approach partitions the data according to primitives by choosing the models such that the description is *best* in terms of global shape and error of fit. Initially seed regions are placed at arbitrary locations in the data and models of each primitive approximated (Fig. 1). Grossly mismatching models may be rejected at this stage. An iterative grow and select phase is then operated. All valid models are grown for an equal number of steps (see left of Fig. 2); note that the models are allowed to overlap. The resulting models are then inspected and some are selected for




Fig. 2. Grown and selected regions at an intermediate stage; final segmentation

further growing. Optimal models from the overlapping sets are selected on the basis of the following criteria:

Area — the number of image elements contained in the model

Error of Fit — the maximum or average distance to the model

Quality of Fit — the number of parameters used to describe the model

Surface Type — the class of surface of the model

When different models have similar goodness of fit, some types of model may be preferred to others, and the **Surface Type** is used to impose an ordering of selection in the way suggested by Besl and Jain [2, 3]. This is, typically, done in terms of increasing surface type complexity (*i.e.* plane, sphere, cylinder, cone, torus), where the simpler surface would be chosen first. A weighted sum of the above criteria is used as a cost-benefit measure to select the optimal model or models.

Models that simply have a poor error of fit (even if not overlapping other models) are also rejected at the selection stage.

The rest of this paper discusses novel fitting methods for spheres, cylinders, cones and tori which are used in this segmentation framework.

3 Surface Fit and General Nonlinear Least-Squares

Initially, we outline some basic concepts for least-squares fitting of simple geometric surfaces, and then review previous approaches to this problem.

Let us assume that each of the three-dimensional points \mathbf{p}_i for $i = 1, \dots, m$ lies close to the same member of a family of surfaces which can be parameterized by $\mathbf{s} \in G \subseteq \mathbb{R}^s$ where G is an open set. Let d be a function which is defined as the distance of the point $\mathbf{p}_i \in \mathbb{R}^3$ from that surface in the family identified by \mathbf{s} . Throughout, d will be called the ‘true’ distance function of the surface (fitting methods generally rely on approximations to this distance as will be explained later).

A surface which goes through all the points can be viewed as that member of the family which corresponds to the solution of the simultaneous system of m equations:

$$d(\mathbf{s}, \mathbf{p}_i) = 0, \quad \text{for } i = 1, \dots, m. \quad (1)$$

Since the number of points m is usually much greater than the number of degrees of freedom s , this system of equations is overdetermined and in general cannot be solved. However it is possible to solve it in the least-squares sense, *i.e.* to find that surface which is the best fit to the points ‘on the average’, minimising

$$\sum_{i=1}^m d(\mathbf{s}, \mathbf{p}_i)^2 \quad (2)$$

Sometimes we may have additional non-linear constraints expressed as

$$H(\mathbf{s}) = \mathbf{0} \in \mathbb{R}^t \quad (3)$$

for $t < s$. For example, the surface form in use might describe a general quadric, but we may wish to impose that the surface found is a cylinder, which can be done by imposing constraints on \mathbf{s} . Using the principle of Lagrangian multipliers to include these constraints, a non-linear generalised eigenvalue problem results, which is not easy to solve. A simpler approach is to use Eqn. 3 to eliminate t unknowns, to reduce the problem to an unconstrained optimisation problem in a lower dimensional space. This is the method we use.

Usually the family of surfaces is defined as points satisfying an implicit equation:

$$f(\mathbf{s}, \mathbf{x}) = 0, \quad \text{for } \mathbf{x} \in \mathbb{R}^3, \quad (4)$$

where \mathbf{s} is the family parameter. Although if we fix \mathbf{s} , f and d have the same roots in space, they may behave quite differently for points which do not lie on the surface. Thus, if instead of Expression 2, one minimises $\sum f^2$ this may give quite different results. However, this approach can be justified if *both* the function f and the constraint H in Eqn. 3 are of particularly simple form. If f is linear and H is quadratic in terms of the parameters then linear generalised eigenvalue techniques work (*e.g.* see [8]). If f is non-linear but H is still quadratic then one can try Taubin’s generalised eigenvector fit [21]. Nevertheless, the choice of form for f influences the behaviour of the non-linear fitting algorithm, and consequently the quality of the solution. For fitting ellipses, Rosin [20] shows that choosing f carelessly can lead to severely biased estimates for \mathbf{s} . Below we give ‘fairly good’ f functions (*i.e.* for which f behaves much like d near the surface) which are highly non-linear, and which have no additional constraints on the parameters \mathbf{s} .

4 Approximating the True Distance

As far as possible one has to avoid singularities of $d(\mathbf{s}, \mathbf{p}_i)$ in the range where solutions may lie. These may be places where some denominator in $d(\mathbf{s}, \mathbf{p}_i)$

vanishes, or where the distance function is not differentiable. Using Euclidean metrics such singularities arise frequently since the Euclidean distance from a given fixed point is *itself* singular in this sense. Nevertheless, most of these nonlinear singularities are only *computational*, *i.e.* inessential discontinuities in the mathematical sense, meaning that a limiting value of the distance function can still be found for the critical parameter value. Even so, the computation of the distance function (or its derivatives) can be unstable at such points, since it may require the subtraction of similar quantities *etc.*

Avoiding the effects of singularities uses various techniques. Firstly, one uses a suitable parameterization where the critical values do not lie on the border of G . Secondly, one changes the definition of $d(\mathbf{s}, \mathbf{p}_i)$ slightly in order to get rid of singularities. We shall say that this modified definition is *faithful* to the true Euclidean distance function if, firstly, the function is 0 where the true distance is 0, and, secondly, at these points, the derivatives with respect to the parameters are the same for the true distance and the modified definition.

Faithful distance functions can be obtained if one approximates square roots within $d(\mathbf{s}, \mathbf{p}_i)$ in the following way. Suppose that the distance function is of the following form:

$$d(\mathbf{s}, \mathbf{p}_i) = \sqrt{g} - h, \quad (5)$$

where both g and h may depend on the parameter vector \mathbf{s} and on the point \mathbf{p}_i in three space. In order to get rid of the square root we might try to minimise $\sum (g - h^2)^2$ instead of Eqn. 2 since $d = 0$ when $g = h^2$. Unfortunately, the effect is now that we are searching for the surface which fits best in terms of the *square* of the distance instead of the distance. This transformation undesirably amplifies the importance of the further points, and flattens the goal function in the neighbourhood of the solution. Instead, let us put

$$\tilde{d} = \frac{g - h^2}{2h} = d + \frac{d^2}{2h}. \quad (6)$$

and let us minimise

$$\sum \tilde{d}^2(\mathbf{s}, \mathbf{p}_i) = \sum \frac{(g(\mathbf{s}, \mathbf{p}_i) - h^2(\mathbf{s}, \mathbf{p}_i))^2}{4h^2(\mathbf{s}, \mathbf{p}_i)}. \quad (7)$$

Under very general assumptions (see [16]), this function is faithful to the Euclidean distance.

In summary, in our approach we start with an exact expression for the distance d . This is replaced by a simplification which is easier to compute, but which still has the same zero set and derivatives at the zero set. In contrast, similar work by Taubin [21] starts from a parameterized family of implicit functions $f = 0$. He notes that while f itself is not a good approximation to d , $f/|\nabla f|$ is much better, *i.e.* he replaces the original implicit function with a new one whose value is a better approximation to d . Although Taubin does not state so explicitly, it is clear that it is better because the derivatives with respect to

spatial parameters of this function are the same as those of the distance function. In practice Taubin's approach can be used only if f is linear with respect to the parameters, and the system to be solved then includes a quadratic constraint. Another difference is that our approach is better behaved with respect to singularities.

5 Fitting Spheres, Cylinders, Cones and Tori

The linear least-squares fitting of second order curves and surfaces has been considered by several authors recently (see [18], [19], [9], [8]). However, specific linear methods still do not exist for right cylinders and cones—the reason is that the equations expressing the conditions for a quadric to be a right cylinder or a cone are not quadratic. If general linear methods are used for algebraic second order surfaces the solutions found are usually not right cylinders or cones and may even be very different from the optimum surfaces of such type. In this sense, algebraic techniques which use the value of the implicit quadratic form as the 'distance' from the surface approximate the true geometric distance in an unfaithful way.

The situation is much simpler for spheres since straightforward algebraic methods work in this case: under a suitable normalisation the minimised algebraic distance will reflect the geometric distance as well. For example, the method in [18] minimises

$$\sum_i (A(x_i^2 + y_i^2 + z_i^2) + Dx_i + Ey_i + Fz_i + G)^2 \quad (8)$$

under the condition

$$D^2 + E^2 + F^2 - 4AG = 1 \quad (9)$$

which is basically equivalent to our minimisation in Eqn. 7. (Note that the simple constraint $A = 1$ may give quite unfaithful results as shown in [18].) Nevertheless, we give our non-linear method for spheres in the next section as an illustration of our method, as it has certain advantages.

Nonlinear methods which take into account the true geometric distance match well with the requirements of our reverse engineering method. Those points belonging to the same surface are selected by means of a segmentation technique which usually provides an initial approximation for the parameters of each surface. Starting from these, at the expense of some computing time, one can obtain a more accurate fit. Our nonlinear methods also work well in other applications where an initial approximate fit for the surface is known.

Earlier nonlinear estimation approaches usually worked with cylinders and spheres. As a rule the equations contain positional parameters of centres or axes and so they become ill-conditioned in limiting situations (see e.g. [4]), which is unacceptable if automatic segmentation is the objective. Our nonlinear methods have been carefully designed to overcome this problem.

5.1 Sphere Fitting

For non-linear least-squares fit the parameterization of the sphere will be the following. Suppose that the *closest* point of the sphere (*not* its centre) to the origin is $\varrho \mathbf{n}$, where $|\mathbf{n}| = 1$ and the radius of the sphere is $1/k$. Then if \mathbf{p} is an arbitrary point in space, the distance of this point from the surface of the sphere is:

$$d(\mathbf{s}, \mathbf{p}) = \left| \mathbf{p} - \left(\varrho + \frac{1}{k} \right) \mathbf{n} \right| - \frac{1}{k} = \sqrt{\langle \mathbf{p} - \left(\varrho + \frac{1}{k} \right) \mathbf{n}, \mathbf{p} - \left(\varrho + \frac{1}{k} \right) \mathbf{n} \rangle} - \frac{1}{k} \quad (10)$$

Since this function is of the form Eqn. 5 we can apply Approximation 6 to give

$$\tilde{d}(\mathbf{s}, \mathbf{p}) = \frac{k}{2} (|\mathbf{p}|^2 - 2\varrho \langle \mathbf{p}, \mathbf{n} \rangle + \varrho^2) + \varrho - \langle \mathbf{p}, \mathbf{n} \rangle, \quad (11)$$

or if one introduces the notation

$$\hat{\mathbf{p}} = \mathbf{p} - \varrho \mathbf{n}, \quad (12)$$

one obtains

$$\tilde{d}(\mathbf{s}, \mathbf{p}) = \frac{k}{2} |\hat{\mathbf{p}}|^2 - \langle \hat{\mathbf{p}}, \mathbf{n} \rangle. \quad (13)$$

Note that here $\hat{\mathbf{p}}$ is the expression of \mathbf{p} with respect to an origin at $\varrho \mathbf{n}$. Now let us parameterize \mathbf{n} using polar coordinates, so in the usual way

$$\mathbf{n} = (\cos \varphi \sin \vartheta, \sin \varphi \sin \vartheta, \cos \vartheta), \quad (14)$$

where ϑ is the angle between \mathbf{n} and the z axis and φ is the angle between the projection of \mathbf{n} onto the plane $z = 0$ and the x axis. Differentiating \mathbf{n} with respect to φ and ϑ one obtains two partial derivative vectors which are orthogonal to each other and to \mathbf{n} (superscripts denote derivatives); these will be used later:

$$\mathbf{n}^\varphi = (-\sin \varphi \sin \vartheta, \cos \varphi \sin \vartheta, 0) \quad (15)$$

$$\mathbf{n}^\vartheta = (\cos \varphi \cos \vartheta, \sin \varphi \cos \vartheta, -\sin \vartheta). \quad (16)$$

Thus \mathbf{n} and hence \tilde{d} can be parameterized without constraints by $\mathbf{s} = (\varrho, \varphi, \vartheta, k)$.

Note that unlike Expression 8, Eqn. 11 is nonlinear but it behaves well as the curvature of the sphere decreases, as in that case, $k \rightarrow 0$, all the terms are bounded, and Eqn. 11 reduces to the expression that would be used for least-squares plane fitting. In contrast, observe that some of the terms will tend to infinity both in the objective function given in Eqn. 8 and in the constraint given in Eqn. 9.

5.2 Right Circular Cylinder Fitting

The parameterization used for the cylinder is similar to that for the sphere. The closest point of the cylinder to the origin is $\varrho \mathbf{n}$, where $|\mathbf{n}| = 1$. Assume that




Fig. 3. Parameterization of the cylinder

the direction of the axis of the cylinder is \mathbf{a} with $|\mathbf{a}| = 1$, and the radius of the cylinder is $1/k$. Note that $\langle \mathbf{n}, \mathbf{a} \rangle = 0$. Now let us suppose that \mathbf{p} is an arbitrary point in space whose distance from the surface of the cylinder is to be found. This is done by finding its distance from the symmetry axis and from that subtracting the radius of the cylinder (see Fig. 3):

$$d(\mathbf{s}, \mathbf{p}) = \left| (\mathbf{p} - (\varrho + \frac{1}{k})\mathbf{n}) \times \mathbf{a} \right| - \frac{1}{k}$$

$$= \sqrt{|\mathbf{p} - (\varrho + \frac{1}{k})\mathbf{n}|^2 - \langle \mathbf{p} - (\varrho + \frac{1}{k})\mathbf{n}, \mathbf{a} \rangle^2} - \frac{1}{k} \quad (17)$$

Since this function is of the form in Eqn. 5 we can apply Approximation 6 to obtain

$$\tilde{d}(\mathbf{s}, \mathbf{p}) = \frac{k}{2} (|\mathbf{p}|^2 - 2\varrho \langle \mathbf{p}, \mathbf{n} \rangle - \langle \mathbf{p}, \mathbf{a} \rangle^2 + \varrho^2) + \varrho - \langle \mathbf{p}, \mathbf{n} \rangle = \frac{k}{2} |\hat{\mathbf{p}} \times \mathbf{a}|^2 - \langle \hat{\mathbf{p}}, \mathbf{n} \rangle, \quad (18)$$

where $\hat{\mathbf{p}} = \mathbf{p} - \varrho \mathbf{n}$ as in Eqn. 12. Using appropriate parameterizations for \mathbf{n} and \mathbf{a} we wish to minimise the function

$$\sum_i \tilde{d}^2(\mathbf{s}, \mathbf{p}_i).$$

Let us make some observations about the right hand side of Eqn. 18. Firstly, it is linear in the curvature k if all other parameters are fixed, which results in a *separable* non-linear least squares problem (see *e.g.* [6]). Such problems can be easier to solve than the fully non-linear case. Clearly, an initial estimate for k is not needed if we have estimates for the other parameters, as an initial value for k can be found by solving a linear least-squares problem in which all other

parameters are fixed. Note that Eqn. 18 behaves well as k gets smaller (ϱ is bounded within sensible limits by the geometric configuration of the scanner); compare Eqn. 17 which subtracts two large quantities as k becomes small. In the limit as $k \rightarrow 0$ we get $\vec{d} = \varrho - \langle \mathbf{p}, \mathbf{n} \rangle$, and as before the problem reduces to linear least-squares fitting of a plane.

Again, we wish to parameterize \mathbf{n} and \mathbf{a} to satisfy the constraints

$$|\mathbf{n}| = |\mathbf{a}| = 1, \quad \langle \mathbf{n}, \mathbf{a} \rangle = 0;$$

again we use polar coordinates. The parameterization for \mathbf{n} was introduced in Eqn. 14, and Eqns. 15 and 16 are the partial derivatives of \mathbf{n} . Thus if we put

$$\overline{\mathbf{n}^\varphi} = (-\sin \varphi, \cos \varphi, 0) = \frac{\mathbf{n}^\varphi}{\sin \vartheta}, \quad (19)$$

then \mathbf{n}^ϑ , $\overline{\mathbf{n}^\varphi}$ and \mathbf{n} form an orthonormal basis. Thus we can parameterize \mathbf{a} as follows:

$$\mathbf{a} = \mathbf{n}^\vartheta \cos \alpha + \overline{\mathbf{n}^\varphi} \sin \alpha, \quad (20)$$

where α is the angle between \mathbf{a} and \mathbf{n}^ϑ . Thus, \mathbf{n} and \mathbf{a} are parameterized through φ , ϑ and α by means of Eqns. 14, 16, 19 and 20.

5.3 Right Circular Cone Fitting

The parameterization used for the cone is similar to that for the cylinder. Let $\varrho \mathbf{n}$ with $|\mathbf{n}| = 1$ be that point on the cone surface for which a line in the direction of the surface normal passes through the origin. (Hence \mathbf{n} is a normal to the cone.) Let the non-zero principal curvature of the cone at the point $\varrho \mathbf{n}$ be k . Let us denote the unit direction of the axis of the cone by \mathbf{a} . \mathbf{n} is parameterized by φ and ϑ as in Eqn. 14. Since \mathbf{n} and \mathbf{a} are not in this case perpendicular, \mathbf{a} can be parameterized freely by two polar coordinate angles, σ and τ

$$\mathbf{a} = (\cos \sigma \sin \tau, \sin \sigma \sin \tau, \cos \tau), \quad (21)$$

where τ is the angle between \mathbf{a} and the z axis and σ is the angle between the projection of \mathbf{a} onto the plane $z = 0$ and the x axis. The six parameters $(\varrho, \varphi, \vartheta, k, \sigma, \tau)$ characterise the right circular cone surface.

In order to understand how this works, let the half angle of the cone be ψ (see Fig. 4), and the position of the apex of the cone be \mathbf{c} ; we shall express ψ and \mathbf{c} using the above parameters later. Let the angle between the axis of the cone and $\mathbf{p} - \mathbf{c}$ be ω . Using these, the distance of \mathbf{p} , an arbitrary point, from the cone surface is given by (note that \mathbf{p} may not lie in the plane of \mathbf{n} and \mathbf{a})

$$d(\mathbf{s}, \mathbf{p}) = |\mathbf{p} - \mathbf{c}| \sin(\omega - \psi) = |\mathbf{p} - \mathbf{c}| \sin \omega \cos \psi - |\mathbf{p} - \mathbf{c}| \cos \omega \sin \psi.$$

Without loss of generality we can suppose that both ψ and ω are acute angles. Since the direction of the axis, \mathbf{a} is a unit vector we have:

$$d(\mathbf{s}, \mathbf{p}) = |(\mathbf{p} - \mathbf{c}) \times \mathbf{a}| \cos \psi - |\langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle| \sin \psi. \quad (22)$$




Fig. 4. Parameterization of the cone

Moreover since the angle between \mathbf{n} and \mathbf{a} is the complementary angle to ψ we have:

$$\cos \psi = |\mathbf{n} \times \mathbf{a}| \quad \sin \psi = |\langle \mathbf{n}, \mathbf{a} \rangle|. \quad (23)$$

Thus from Eqn. 22 one obtains:

$$\begin{aligned} d(\mathbf{s}, \mathbf{p}) &= |(\mathbf{p} - \mathbf{c}) \times \mathbf{a}| |\mathbf{n} \times \mathbf{a}| - |(\mathbf{p} - \mathbf{c}, \mathbf{a}) \langle \mathbf{n}, \mathbf{a} \rangle| \\ &= |\mathbf{n} \times \mathbf{a}| \sqrt{|\mathbf{p} - \mathbf{c}|^2 - \langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle^2} - |\langle \mathbf{n}, \mathbf{a} \rangle \langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle|. \end{aligned} \quad (24)$$

Again this function is of the form given in Eqn. 5, and using Approximation 6 gives

$$\tilde{d}(\mathbf{s}, \mathbf{p}) = \frac{|\mathbf{p} - \mathbf{c}|^2 \cos^2 \psi - \langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle^2}{2 \langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle \sin \psi} = \frac{|\mathbf{p} - \mathbf{c}|^2 |\mathbf{n} \times \mathbf{a}|^2 - \langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle^2}{2 \langle \mathbf{p} - \mathbf{c}, \mathbf{a} \rangle \langle \mathbf{n}, \mathbf{a} \rangle}. \quad (25)$$

We now express the position of the apex \mathbf{c} in terms of the normal vector $\mathbf{n}(\varphi, \vartheta)$, the distance ϱ , the curvature k and the axis of the cone $\mathbf{a}(\sigma, \tau)$:

$$\mathbf{c} = \left(\varrho + \frac{1}{k}\right) \mathbf{n} + \gamma \mathbf{a}.$$

Now $\langle \mathbf{c}, \mathbf{n} \rangle = \varrho$, so $\gamma = -1/(k \langle \mathbf{n}, \mathbf{a} \rangle)$, and thus

$$\mathbf{c} = \left(\varrho + \frac{1}{k}\right) \mathbf{n} - \frac{\mathbf{a}}{k \langle \mathbf{n}, \mathbf{a} \rangle}. \quad (26)$$

Substituting this into Eqn. 25 one obtains (again $\hat{\mathbf{p}} = \mathbf{p} - \varrho \mathbf{n}$)

$$\tilde{d} = \frac{|\mathbf{n} \times \mathbf{a}|^2 \left(|\hat{\mathbf{p}} - \mathbf{n}/k|^2 - \langle \hat{\mathbf{p}} - \mathbf{n}/k, \mathbf{a} \rangle^2 \right) - \left(\langle \hat{\mathbf{p}} - \mathbf{n}/k, \mathbf{a} \rangle \langle \mathbf{n}, \mathbf{a} \rangle + 1/k \right)^2}{2 \left(\langle \hat{\mathbf{p}} - \mathbf{n}/k, \mathbf{a} \rangle \langle \mathbf{n}, \mathbf{a} \rangle + 1/k \right)},$$

Using Pythagoras' theorem it is easy to see that the coefficient of $1/k^2$ in the numerator is zero. Multiplying both the numerator and the denominator by k we arrive at

$$\tilde{d}(\mathbf{s}, \mathbf{p}) = \frac{\frac{k}{2}(|\mathbf{n} \times \mathbf{a}|^2 |\hat{\mathbf{p}}|^2 - \langle \hat{\mathbf{p}}, \mathbf{a} \rangle^2) - \langle \hat{\mathbf{p}}, \mathbf{n} \rangle |\mathbf{n} \times \mathbf{a}|^2}{k \langle \hat{\mathbf{p}}, \mathbf{a} \rangle \langle \mathbf{n}, \mathbf{a} \rangle + |\mathbf{n} \times \mathbf{a}|^2}. \quad (27)$$

5.4 Torus Fitting

Our approach for tori is again similar. A torus can be parameterized through seven unconstrained parameters. A torus can be obtained as a surface swept by a circular disc rotated around an axis in the plane of the circle. We call the radius of the disc the 'minor radius' and the distance of the centre of the disc from the axis the 'major radius' of the torus. Tori whose major radius is *smaller* than the minor one can also be considered. In this case the resulting surface is self-intersecting, and it is necessary to distinguish the different parts in the following. The smaller arcs sweep a 'lemon-torus' (*i.e.* the inner part of the torus surface), while the larger arcs sweep an 'apple-torus' (*i.e.* the outer part of the torus surface); we will also refer to a non-self-intersecting torus as 'apple-shaped'. In special cases the torus may degenerate into a sphere, as the major radius vanishes, or into a cone, as the minor radius tends to infinity. Our equations appropriately reduce to those for sphere or cone fitting. (If the major radius tends to infinity the torus becomes a cylinder. This case will be singular, but mathematically will be close to the cylinder fit.) The parameterization used for the torus is the following. As before, the point on the torus where a line through the surface normal passes through the origin is $\varrho \mathbf{n}$, where $|\mathbf{n}| = 1$. The principal curvature of the torus corresponding to the minor radius at the point $\varrho \mathbf{n}$ is k (*i.e.* the radius of the disk is $1/k$). The other principal curvature is s , and the corresponding centre of curvature lies on the axis of the torus. Let the unit direction vector of the torus axis be \mathbf{a} . (See Fig. 5.) We parameterize \mathbf{n} by φ and ϑ as in Eqn. 14 and the unit vector \mathbf{a} as in Eqn. 21:

$$\mathbf{a} = (\cos \sigma \sin \tau, \sin \sigma \sin \tau, \cos \tau).$$

The unconstrained parameters $(\varrho, \varphi, \vartheta, k, s, \sigma, \tau)$ entirely characterise the torus surface.

We may now state (the proof may be found in [16]) that a non-linear distance function for tori which is faithful up to the first derivative is

$$\tilde{d}(\mathbf{s}, \mathbf{p}) = \tilde{d}_0(\varrho, \varphi, \vartheta, k, \mathbf{p}) - \delta_\epsilon(\varrho, \varphi, \vartheta, k, s, \sigma, \tau, \mathbf{p}) \quad (28)$$

where \tilde{d}_0 is the approximate distance function for the sphere given in Eqn. 11:

$$\tilde{d}_0 = \frac{k}{2} (|\mathbf{p}|^2 - 2\varrho \langle \mathbf{p}, \mathbf{n} \rangle + \varrho^2) + \varrho - \langle \mathbf{p}, \mathbf{n} \rangle = \frac{k}{2} |\hat{\mathbf{p}}|^2 - \langle \hat{\mathbf{p}}, \mathbf{n} \rangle,$$

while

$$\delta_\epsilon = \left(\frac{k}{s} - 1 \right) \left[\epsilon \cdot \text{sign} \left(\frac{k^2}{s} - k \right) |(\hat{\mathbf{p}} - \mathbf{n}/s) \times \mathbf{a}| |\mathbf{n} \times \mathbf{a}| + \langle (\hat{\mathbf{p}} - \mathbf{n}/s) \times \mathbf{a}, \mathbf{n} \times \mathbf{a} \rangle \right] \quad (29)$$




Fig. 5. Parameterization of the torus

where $\epsilon = +1$ for an apple torus surface and $\epsilon = -1$ for a lemon torus surface.

This function behaves well in limiting situations. If $k = s$ then from Eqn. 29 we simply get the distance expression for a sphere given in Eqn. 13. If $k \rightarrow 0$ and s is bounded from below then either for $\epsilon = -1$ or for $\epsilon = +1$ we get the distance expression for a cone given in Eqn. 24. If $s \rightarrow 0$ then Eqn. 28 degenerates to the distance function of a cylinder given in Eqn. 18 with axis direction $\mathbf{n} \times \mathbf{a}$. (See [16].)

6 Initial Estimates

To find the solution of any of the above nonlinear least-squares problems, an iterative technique is used; in practice we use the Levenberg-Marquardt method (see [6]). Any such algorithm requires some good initial estimate of the solution which is then refined. Here we give one method of finding such initial estimates.

The first step of this process is to find an estimate of the rotational axis (except for sphere fitting). A method is given to do this in [16] which is based only on estimates of the surface normal vector. It computes the axis from a number of four-tuples of normals. An alternative approach which finds the best (least-squares) rotational axis has recently been suggested by Pottmann [17].

We then pick a surface point at which we have a normal vector estimate and locate the origin at this point, which we call the 'base point'. Thus we have an estimate for \mathbf{n} and hence φ, ϑ ; the initial estimate for ϱ is 0. (Note that the solution surface need not pass through the base point, since ϱ can change.)

For sphere fitting, these values alone are sufficient, and an estimate for k can now be found by solving a linear least-squares problem as mentioned earlier.

For cylinder fitting, we adjust the normal to be perpendicular to the axis. We then compute α as the angle between \mathbf{a} and \mathbf{n}^ϑ and compute the distance of the base point from the axis, which is $1/k$ signed with the direction of \mathbf{n} .




Fig. 6. Simulated part depth-map; seed placement




Fig. 7. Simulated part intermediate grown regions; final segmentation

In the case of cone fitting, after estimating the rotational axis \mathbf{a} we compute the distance of the base point from the axis *along* the estimated normal line in order to obtain $1/k$.

In the torus case, for better conditioning one has to pick a base point at which the normal subtends as large an angle as possible with the estimated axis. We get $1/s$ as the distance along the normal between the base point and the estimated axis. For simplicity, we can put $k = 0$ (*i.e.* start from a cone) and try fitting both $\epsilon = +1$ (apple-torus) and $\epsilon = -1$ (lemon-torus). More robustly, one can opt to estimate principal curvatures of the surface at the base point. As s is one of the principal curvatures we can compute the other, k , even if we only estimate the Gaussian curvature. We can then determine ϵ by noting on which sheet of the torus the base point lies. Thus $\epsilon = +1$ if $|k| > |s|$ or $ks < 0$ and $\epsilon = -1$ otherwise (the decision should be clear provided we have chosen a well placed base point as described above). We assume here that the point set being fitted does not contain points belonging to both the apple and lemon sheets of the same torus simultaneously. This is very unlikely to happen in practice, but if it is considered to be a possibility, before fitting we should separate the points into two sets, one for each sheet, using curvature estimates and the given criterion.




Fig. 8. Heriot-Watt part depth-map; seed placement




Fig. 9. Heriot-Watt part intermediate grown regions; final segmentation

7 Results and Conclusion

The fitting routines described in this paper have been tested using the segmentation approach outlined earlier. In the case of simulated 3D point data, which was accurate to five significant digits, in all cases the models recovered fitted the data to an accuracy of at least four significant digits, and segmentations consistent with the underlying geometric primitives were obtained. An example of the final segmentation obtained for the peg composed of a cylinder and sphere shown in Fig. 1 is given in Fig. 2. A further example containing a cone, a cylinder and two spheres is shown in Figs. 6 and 7.

Our fitting methods also work well in practice with real scanner data, which is somewhat less accurate than the simulated data. The segmentation process for a depth-map of the well-known ‘Heriot-Watt’ test part captured using a 3D Scanners *Replica* device is illustrated in Figs. 8 and 9. The segmentation results here are not as clean as for the simulated parts, but are adequate for input to the further model building processes of our reverse engineering system.

Whilst our motivation is the reverse engineering of boundary representation solid models from three-dimensional depth-maps of scanned objects, we believe that the fitting methods described in this paper are of interest to the computer vision and CAD communities in general.

In summary, we have described novel methods for the least-squares fitting of spheres, cylinders, cones and tori to point data. We have outlined how they can be used in a segmentation strategy that is capable of extracting these surfaces from three-dimensional data. Initial results show that the accuracy achieved by these methods is good, although space limitations preclude a full presentation here. Our fitting methods have the major advantage of being robust in the sense that as the principal curvatures of the surfaces being fitted decrease (or as they become more equal), the results which are returned naturally become closer and closer to surfaces of 'simpler type', *i.e.* planes, cylinders, or cones (or spheres, in the case of equal curvatures) which best describe the data. Furthermore, our methods inherently avoid all singularities (except in the case of the torus, for which the problem can readily be overcome by choosing the origin appropriately).

References

1. R. Bajcsy, F. Solina, and A. Gupta, *Segmentation versus object representation—are they separable?*, In Analysis and Interpretation of Range Images, Eds. R. Jain and A. K. Jain, Springer-Verlag, New York, 1990.
2. P. J. Besl, *Surfaces in Range Image Understanding*, Springer-Verlag, New York, USA, 1988.
3. P. J. Besl and R. K. Jain, *Segmentation Through Variable-Order Surface Fitting*, IEEE Transactions on Pattern Analysis and Machine Intelligence, **10** (2), 167–192, 1988.
4. R. M. Bolle and D. B. Cooper, *On optimally combining pieces of information, with application to estimating 3-D complex-object position from range data*, IEEE Transactions on Pattern Analysis and Machine Intelligence **8**, 5, 619–638, 1986.
5. R. M. Bolle and B. C. Vemuri, *On Three-Dimensional Surface Reconstruction Methods*, IEEE Transactions on Pattern Analysis and Machine Intelligence, **13** (1), 1–13, 1991.
6. Å. Björk, *Numerical methods for least squares problems*, SIAM. Society for Industrial and Applied Mathematics, Philadelphia, 1996.
7. O. D. Faugeras, M. Hebert, and E. Pauchon, *Segmentation of Range Data into Planar and Quadric Patches*, Proceedings of Third Computer Vision and Pattern Recognition Conference, Arlington, VA, 8–13, 1983.
8. A.W. Fitzgibbon, M. Pilu, and R.B. Fisher, *Direct least-square fitting of ellipses*, 13th International Conference on Pattern Recognition (Washington, Brussels, Tokyo), IAPR, IEEE Computer Society Press, June 1996, Proceedings of the 13th ICPR Conference, Vienna Austria, August 1996.
9. W. Gander, G.H. Golub, and R. Strelbel, *Least-squares fitting of circles and ellipses*, BIT **34**, 558–578, 1994.
10. M. Hebert and J. Ponce, *A new method for segmenting 3-D scenes into primitives*, 6th International Conference on Pattern Recognition (Munich), 836–838, IAPR DAGM, IEEE Computer Society Press, October 1982, Proceedings of the 6th ICPR Conference, Munich Germany, Oct. 19–22, 1982.
11. A. Jaklic, A. Leonardis, and F. Solina. *Segmentor: An object-oriented framework for image segmentation*. Technical Report LRV-96-2, Computer Vision Laboratory, University of Ljubljana, Faculty of Computer and Information Science, 1996.

12. A. Leonardis. *Image analysis using parametric models: model-recovery and model-selection paradigm*, PhD dissertation, University of Ljubljana, Faculty of Electrical Engineering and Computer Science, May 1993.
13. A. Leonardis, A. Gupta, and R. Bajcsy. *Segmentation as the search for the best description of the image in terms of primitives*. Proceedings of the Third International Conference of Computer Vision, Osaka, Japan, 1990.
14. A. Leonardis, A. Gupta, and R. Bajcsy. *Segmentation of range images as the search for geometric parametric models*. International Journal of Computer Vision, **14**, 253–277, 1995.
15. P. Liong, and J. S. Todhunter, *Representation and recognition of surface shapes in range images: a differential geometry approach*, Computer Vision, Graphics and Image Processing, **52**, 1, 78–109, 1990.
16. G. Lukács, A. D. Marshall, and R. R. Martin, *Geometric least-squares fitting of spheres, cylinders, cones and tori*, RECCAD, Deliverable Document 2 and 3, COPERNICUS project, No 1068 (Budapest) (R. R. Martin and T. Várady, eds.), Geometric Modelling Laboratory Studies/1997/5, Computer and Automation Research Institute, Budapest, July 1997.
17. H. Pottmann and T. Randrup, *Rotational and helical surface approximation for reverse engineering*, Tech. Report 43, Institut für Geometrie, Technische Universität Wien, A-1040 Wien, Wiedner Hauptstraße 8–10, Austria, June 1997, Submitted to Computing.
18. V. Pratt, *Direct least-squares fitting of algebraic surfaces*, COMPUTER GRAPHICS Proceedings, vol. 21, Annual Conference Series, no. 4, ACM, Addison Wesley, July 1987, Proceedings of the SIGGRAPH 87 Conference, 145–152, Anaheim, California, 27–31 July 1987.
19. P. L. Rosin, *A note on the least squares fitting of ellipses*, Pattern Recognition Letters **14**, 799–808, 1993.
20. P. L. Rosin, *Analysing error of fit functions for ellipses*, Pattern Recognition Letters **17**, 1461–1470, 1996.
21. G. Taubin, *Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation*, IEEE Transactions on Pattern Analysis and Machine Intelligence **13** (11), 1115–1138, 1991.