

Algorithms for reverse engineering boundary representation models

Pál Benkő^{a,*}, Ralph R. Martin^b, Tamás Várady^a

^aComputer and Automation Research Institute, Budapest, Hungary

^bCardiff University, UK

Accepted 15 January 2001

Abstract

A procedure for reconstructing solid models of conventional engineering objects from a multiple-view, 3D point cloud is described. (Conventional means bounded by simple analytical surfaces, swept surfaces and blends.) Emphasis is put on producing accurate and topologically consistent boundary representation models, ready to be used in computer aided design and manufacture. The basic phases of our approach to reverse engineering are summarised, and related computational difficulties are analysed.

Four key algorithmic components are presented in more detail: efficiently segmenting point data into regions; creating linear extrusions and surfaces of revolution with smooth, constrained profiles; creating the topology of B-rep models; and finally adding blends. The application of these algorithms in an integrated system is illustrated by means of various examples, including a well-known reverse engineering benchmark. © 2001 Elsevier Science Ltd. All rights reserved.

1. Introduction

Reconstructing objects has been a central problem in computer vision for a long time [7,8,22]. Technological advances in laser scanning now make it possible to produce multiple 3D point clouds with high density and high accuracy, and in this way it has become a realistic expectation to generate accurate models with correct topology which can be directly utilised by CAD/CAM systems. There is a wide diversity of reverse engineering methods for converting measured data into geometric models [38]. Methods differ in the quality of measurements, the quantity of points, the geometric characteristics of the objects in question, the amount of user interaction required, and the final representation used for the model.

About 20 years ago, the two main areas of geometric modelling development separately addressed (i) geometrically demanding free-form surfaces and (ii) topologically complex solid models with simple surface elements. These techniques were eventually integrated into commercial CAD/CAM systems. The current situation in reverse engineering is somewhat similar. For free-form shapes, the current state-of-the-art relies on techniques in which the user manually segments point clouds. Segmentation is followed by fitting high-quality surfaces [24,25] and the result is typically a collection of trimmed parametric surface

patches. For conventional engineering objects (i.e. those bounded by simple analytic surfaces, swept surfaces and blends), the major efforts are directed towards *automatic methods* of building accurate and consistent CAD models with little or no user interaction. These approaches rely on various assumptions concerning the regularity of the objects, which make it possible to automate the process.

Although there are many valuable contributions in computer graphics and computer vision concerning object reconstruction, the majority of these cannot directly be applied to building CAD/CAM models for mechanical engineering. Some produce fair, triangulated, approximations suitable for graphics (see e.g. [23]), while others represent the boundary of objects using special surface representations such as A-patches or superquadrics [2,31]. Such representations are of limited use as commercial CAD/CAM systems must use standard representations for interoperability.

Approaches exist which use standard free-form surfaces, such as B-splines, to reconstruct conventional engineering objects [13,17,26,37]. Two difficulties arise. If multiple surfaces are used, it is hard to find good segmenting curves, and to assure C^0 or higher continuity between the adjacent surfaces. On the other hand, fitting a single B-spline surface to the whole data cannot properly reproduce the piecewise nature of the surface with appropriate curvature distribution. Even if these limitations could be overcome, another problem remains. Representing the separate analytical surfaces of the original design is important in order to be

* Corresponding author.

E-mail address: benko@sztaki.hu (P. Benkő).

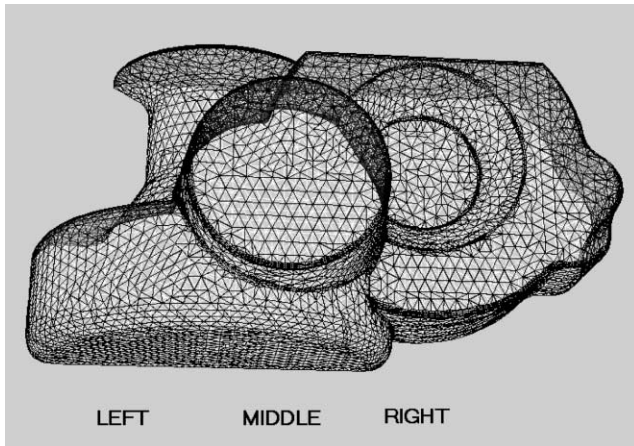


Fig. 1. Triangulated and decimated data points.

able to edit the model in a meaningful way, to be able to plan correct manufacturing operations, and for other related purposes. A free-form model does not provide this information.

The above arguments justify dedicated approaches for conventional engineering objects, such as the well-known benchmark in Fig. 1, which was put forward a few years ago by Hoschek [24]. From a functional point of view, this part is the union of three components. LEFT has rotational symmetry, MIDDLE is a part of a blended cone, and RIGHT is bounded by two planar faces and a cone, as well as a sequence of extruded side faces. All intersection edges have been blended by constant radius rolling ball blends. Clearly, representing this shape with a single B-spline surface would be inadequate, even if possible.

We now state the problem we wish to solve. We assume an ungridded, dense set of point data is available, formed by merging multiple views. We assume the data come from a conventional object bounded by relatively large *primary surfaces* of simple analytic surface type (planes, natural quadrics and tori), which must be accurately represented as such, not as B-splines. Linear extrusions (formed by sweeping a profile curve along a straight line perpendicular to it) and surfaces of revolution must be recognised and represented by means of swept profiles. Fixed-radius *blends* must also be reconstructed as small transition surfaces which smoothly join the primary surfaces. If two adjacent faces of the original object are tangent continuous, the reconstructed model should also guarantee smoothness. Since surface fitting to measured datasets is inherently inaccurate, we must recognise appropriate engineering constraints such as perpendicularity, parallelism, and concentricity, and enforce them on the model either at the fitting stage or as a postprocess.

In comparison, the majority of previous reverse engineering approaches only considered single view range data, and only concerned itself with simple surfaces [10,16,21,34].

In Section 2, we outline our automatic reverse engineering approach for conventional objects, surveying relevant

ideas and identifying the successive computational steps. Then, in Sections 3–6, four particular algorithmic components are chosen and discussed in more detail: segmenting point clouds, reconstructing linear extrusion surfaces and surfaces of revolution, building the topology of B-rep models, and adding blends.

2. Basic phases of reconstructing solid models

Our methodology for the reverse engineering procedure can be decomposed into the following steps:

1. data capture;
2. merging multiple point clouds;
3. triangulation/decimation;
4. segmentation;
5. surface fitting for simple surfaces;
6. reconstructing linear extrusions and surfaces of revolution;
7. reconstructing smooth multiple regions;
8. building an adjacency graph;
9. constraint identification, constrained fitting;
10. B-rep model creation, further beautification;
11. blend reconstruction.

For *data capture*, typically laser scanners are used, which produce measured datasets with high density and reasonable accuracy, although outliers may occur, particularly near silhouette curves and concave edges. There are missing portions due to occlusion in single views, meaning that discontinuities in depth occur, and this may also hold even for merged point clouds.

A single view represents only part of an object. In order to get a complete (or at least a multiple-view) model, point clouds need to be *merged*, i.e. appropriate transformations need to be found by means of which the different datasets can be represented in a common coordinate system. This is often called *registration*, which is an important problem in computer vision [5,12]. Iterative closest point methods match nearest points in several datasets. Such methods tend to be slow, and require relatively good initial estimates of the registration. Some progress has been made towards overcoming these difficulties [30]. An efficient hierarchical method can be used based on computing characteristic curves in the point data and then special points on those curves. Registration is first done using the points, then refined using the curves and finally made accurate using the whole point set. This approach is even effective for obtaining an initial registration. Secondly, iterative *reciprocal* closest point methods are more robust when there is relatively little overlap between the datasets.

Other registration methods use calibrating balls or special markers to find the correspondence between the individual point sets. High-precision rotary tables are often used, since transformations can robustly be computed; however, in this

case the bottom face and other downward pointing faces are not visible.

In fact, there is more than one approach for producing models from multiple views. If it is done by *merging point clouds* as described above, a vast data structure is obtained with large, united point regions, which may suffer from registration inaccuracies and uneven, irregular point distributions. The alternative is to create single view, partial B-rep models; in this case we have to deal with discontinuities, but simpler models are obtained. However, the single view models must be consistent and accurate to be able to *combine* them using Boolean operations to produce a final model. The pros and cons of point cloud merging and combining single-view models are discussed in detail in Ref. [3].

Even a single point cloud may contain at least 10^4 and perhaps as many as 10^8 data points. In order to determine the surface of an object, first the neighbourhoods and connectivity of the data points must be computed. This is a difficult problem, particularly if we have a merged dataset with noise, uneven point distributions, discontinuities, holes and disjoint parts. There are various techniques to build a consistent *triangulation*, which differ in their basic assumptions, their generality and their computational efficiency [2,14,23]. Methods based on variants of α -shapes are popular, but require careful implementation and the proper selection of parameters to achieve successful results [18]. As an alternative, we use an efficient method suggested by Kós [28], in which locally created Delaunay triangulations are merged in a consistent manner.

For efficiency reasons, the original triangulation often needs to be *decimated*, forming a so-called *carrier surface*. Decimation needs to be an adaptive process taking into account the local surface geometry; it is a widely studied problem to control the amount of data displayed in computer graphics [20]. The carrier surface reflects the original topology of the point cloud; each data point is either a vertex of the triangulation or associated with the closest triangle of the carrier. The carrier surface plays an important role in speeding up further computations.

One of the most crucial elements of reverse engineering is *segmentation*, i.e. identifying point regions which are disjoint subsets and mark out connected portions within the whole point cloud sharing some common property. These are considered to be pre-images of single faces (or possibly multiple faces) of the B-rep model. For each face a single surface is *fitted* to the related points within a given tolerance. Various methods for fitting simple analytic surfaces have been described [8,15,16,19,32,40]. We believe that the use of geometric, or geometrically faithful, distances is more satisfactory for surface reconstruction than algebraic distances.

In general, segmentation is a complex process [38]. Often iterative *region growing* techniques are applied [7,31,37]. While it is possible to apply the region growing paradigm to segment conventional engineering objects [33], in this paper

we prefer another solution called *direct segmentation* [40]. This method is an efficient, non-iterative approach, which satisfies the requirements we set for the accurate reconstruction of solid models.

Direct segmentation is based on the fact that it is possible to compute local characteristic quantities (e.g. normal direction) within the interior of faces. If these quantities are consistent within a neighbourhood, this implies that we are in the interior of some region. If they change rapidly or fluctuate, we may deduce that we are in between two or more regions. Based on this principle, we can segment the point cloud. We discard points in the neighbourhood of unstable estimates and use only the stable data points for surface fitting. As described in Section 3, direct segmentation first splits the object into smaller regions along sharp edges and small radius blends. For each region an attempt is made to approximate the points by a *single* analytic surface. The regions for which this cannot be done are considered *multiple, smooth regions* and need to be further subdivided. We first search for *translational or rotational symmetries*, and if either is found we determine a smooth, sweeping *profile* (see Section 4) and split the region accordingly. If further multiple, smooth regions remain, we need more sophisticated segmentation methods. These regions contain surface elements with tangential contact and surfaces must be simultaneously fitted with smoothness constraints.

Direct segmentation produces disjoint regions, each of which is approximated by a simple analytic or swept surface. During this process, we have also identified *sharp edges*, which can be computed by surface–surface intersection, and *smooth edges* where surface–surface intersection must not be used. Smooth edges are directly computed by means of *constrained fitting*.

Based on the segmented regions, a so-called *region adjacency graph* is built, which reflects the complete topology and serves as the basis for building the final B-rep model, as discussed in Section 5. The actual construction of B-rep models is generally performed using a solid modelling kernel. We prefer to build so-called *stitched* models, where the individual faces are glued together along their common edges, and so a composite shell is formed gradually. Alternatively, as mentioned earlier, it is possible to produce *volumetric* models by combining prismatic objects constructed from single views [3].

Smooth edges which are the boundaries of *blending surfaces* occur in a different way. Blends are temporarily ignored during segmentation and fitting, and are reconstructed only in the very last phase of model building [27]. Blending surfaces are constructed to be tangential to the primary surfaces: see Section 6.

A crucial step is the *beautification* of the final model [42,43]. This can take place both before and after B-rep model building. In the presence of noisy measured data, the generated model is likely to be imperfect. The exclusion of very small edges and facets, filling of little holes, ensuring orthogonality, parallelism, concentricity, etc. are all

important requirements for real-life CAD/CAM models. The latter types of beautification can also be solved by constrained fitting. As discussed earlier, for composite regions, smoothness constraints are added and the surface elements are simultaneously fitted. Similarly, for example, if two planes are found to be nearly perpendicular, the points of the two related surfaces can be refitted under a perpendicularity constraint, which will lead to a better model. There are many papers on constrained geometric modelling [9], but its relevance to reverse engineering has been recognised only recently [4,41].

Having concluded the overview of the reconstruction process, we continue by providing more details of particular algorithmic phases within the above general framework.

3. Direct segmentation

In this section we describe the basic concept of direct segmentation, but lack of space prevents us from giving full details, for which see Refs. [4,40]. Our purpose is to subdivide the point data into disjoint subsets, which are smooth and to which one (or more) surfaces can locally be fitted. As noted earlier, the fundamental structure of the object is determined by primary surfaces. Along the intersection of two primary surfaces there is either a sharp edge or a small-radius blend. Alternatively, the surfaces may be smoothly connected.

In the first stage of direct segmentation we attempt to separate regions by removing triangles which are in the vicinity of sharp edges or small blends. Thus, disjoint regions are obtained, which can be classified either as *simple*, i.e. which can be approximated by a single analytic surface, or as *multiple*, which need to be further subdivided by smooth edges.

An important feature of the current algorithm is that, for each data point, we estimate a surface *normal vector* and a value which characterises the *planarity* of the point neighbourhood. Our experience shows that the best normal vector estimates are based on fitting a second order algebraic surface to surrounding points in the neighbourhood. This means that first we translate all the points in the neighbour-

hood so that the given point \mathbf{p} goes to the origin (the image point set is denoted by P'), then we look for a symmetric matrix A and unit vector \mathbf{d} which minimise

$$\sum_{\mathbf{x} \in P'} (\mathbf{x}^T A \mathbf{x} + \langle \mathbf{d}, \mathbf{x} \rangle)^2.$$

This leads to a three-dimensional eigenvalue problem and the normal we seek is just \mathbf{d} . $\langle \cdot, \cdot \rangle$ denotes the scalar product of two vectors.

The Planarity is characterised by taking a point and the related least square plane fitted to its neighbours, and computing the normalised fitting errors over a small disk: for relatively smooth primary surfaces this value will be small, while for highly curved parts this value will be greater. There are other alternatives to identify highly curved parts using the triangulation; for example, the angular variation of normals was suggested in Ref. [26], or curvature estimates can be computed [1].

After choosing an appropriate threshold value, planarity filtering will remove those triangles which are classified as belonging to highly curved regions of the point cloud. For example, in Fig. 2, we have filtered out the small blends between the components LEFT, MIDDLE and RIGHT, and the small blends between the side faces and the top and bottom planar faces. The larger blends contained in the linear extrusion and surface-of-revolution regions have been kept. We obtained two smooth, multiple regions, as well as many simple regions (Fig. 3).

Note that our example object is a variant of the original benchmark of Hoschek, which did not contain sufficiently dense data for our purposes, making it impossible to reconstruct blends reliably. Our variant was redesigned and sampled in such a way that we can demonstrate important features of our algorithm.

Given a region, we do not immediately know whether it is simple or multiple. A sequence of hypotheses is tested to decide whether the region's point set can be approximated by some simple surface. If an hypothesis fails, we proceed to try more complex surface types. Finally, multiple regions are decomposed. We always proceed in order of increasing complexity. Fortunately, as simpler surfaces (i) occur more




Fig. 2. Planarity filtering—simple and smooth, multiple regions.




Fig. 3. Segmented object, simple linear extrusions and surfaces of revolution detected.

frequently in mechanical engineering, and (ii) can be detected in a more reliable and efficient manner. The decision procedure thus follows the simplified scheme below.

1. Test whether the region is a plane or a sphere; if so: determine parameters. OTHERWISE
2. Test for translational symmetry; if so:
 - 2.1. Test whether it is a cylinder; if so: determine parameters, ELSE
 - 2.2. Test whether the sweeping contour can be decomposed into a smooth sequence of straight line segments and circular arcs; if so: multiple region, ELSE
 - 2.3. Compute a free-form profile. OTHERWISE
3. Test whether it is a cone; if so: determine parameters. OTHERWISE
4. Test for rotational symmetry; if so:
 - 4.1. Test whether it is a torus; if so: determine parameters. ELSE
 - 4.2. Test whether the sweeping contour can be decomposed into a smooth sequence of straight line segments and circular arcs; if so: multiple region, ELSE
 - 4.3. Compute a free-form profile. OTHERWISE
5. Decompose non-regular, multiple smooth regions
 - 5.1. Detect planes
 - 5.2. Detect linear extrusion sub-regions
 - 5.3. Search for independent axes of rotation, and for each axis detect surface-of-revolution sub-regions

(For 5.2 and 5.3 proceed as in 2 and 4.)

The above tests are based on consecutive least squares procedures for the geometric parameters of these simple surfaces [40]. The scheme shows the significance of detecting translational and rotational symmetries, which are based on the surface normal estimates associated with each point.

We will give details of certain key steps shortly, but in the meantime two minor points in the above procedure are worthy of note. Firstly, consider cone fitting. True least squares cone fitting, of course, leads to a nonlinear system of equations. We may detect a cone as a surface of revolution with a linear profile, but our experiments show that estimating the apex by a least squares fit using the tangent planes at each data point, and computing a least squares fit plane for the normal vectors on the Gaussian sphere provides more reliable results. This process gives the axis direction and the half angle of the cone.

Secondly, consider general free-form surfaces. The last paragraph of the decision hierarchy above should include an additional item (5.4), to extend the concept to also include free-form surfaces. Here, our hypothesis would be that for all isolated regions which have not been confirmed earlier to be a simple analytic or swept surface with either straight-circular or free-form profile, there exists a B-spline surface that can well approximate the remaining data points. Of

course, special care is required in these cases to enforce at least numerical positional and tangential continuity with the surrounding analytic geometric elements.

3.1. Determining the translational direction

We can estimate the translational direction by means of a least squares fit. The normal vectors of a linear extrusion surface must be perpendicular to a common translational direction. We may minimise the deviation of these angles from $\pi/2$, but the simplest method is to minimise the cosine of the angles, which is just the scalar product of the given normal vectors and the unknown direction. The translational characteristic of a shape can be visualised with the help of the Gaussian sphere: the normal vectors of a linear extrusion surface lie on a great circle, that is, in a plane which goes through the origin and is perpendicular to the translational direction.

Formally: if the estimated normal vectors of the region are denoted by $\{\mathbf{n}_i\}$, then we look for the unit vector \mathbf{d} which minimises

$$\sum \langle \mathbf{n}_i, \mathbf{d} \rangle^2.$$

This is a well-known three-dimensional eigenvalue problem.

3.2. Determining the rotational axis

The normal lines of a surface of revolution must intersect a common axis (in a projective sense, i.e. the lines may also be parallel to the axis). The first problem is how to measure the error of a noisy normal line. Using the distance between the axis and the normal line has two problems: firstly, a normal line parallel to the axis does not have zero error; secondly, the error in a line near to the axis is penalised less than a line farther from the axis with the same angular deviation. Another idea is to use the angle between the normal line, and the plane spanned by the axis and the data point corresponding to the normal line. Now, we do not face the first problem as above, but have the opposite of the second problem: the noise in a data point near to the axis is penalised more than in a further point. Using the product of the distance and the sine of the angle defined above, however, eliminates the problems (and has the advantage of being easy to compute as will be shown below). This measure was suggested by Pottmann and Randrup [36].

We cannot go into details here, and just the basic concepts are presented. Lines are represented using Plücker coordinates. A line with direction \mathbf{d}_i through point \mathbf{p}_i is represented by a sextuple $(\mathbf{d}_i, \mathbf{d}_i \times \mathbf{p}_i)$ (the latter term, the *momentum vector* is denoted by $\bar{\mathbf{d}}_i$. Note that it is independent of the choice of \mathbf{p}_i on the line.) The axis a is also given in the same form as $(\mathbf{a}, \bar{\mathbf{a}})$.

Using this formalism, our error function is linear in the coordinates of the unknown axis; we minimise

$$\sum [\langle \bar{\mathbf{d}}_i, \mathbf{a} \rangle + \langle \mathbf{d}_i, \bar{\mathbf{a}} \rangle]^2.$$

We solve the above least squares expression for $(\mathbf{a}, \bar{\mathbf{a}})$ under the constraints $\|\mathbf{a}\| = 1, \langle \mathbf{a}, \bar{\mathbf{a}} \rangle = 0$ (which are needed for normalisation and for the six numbers to be correct Plücker coordinates). Initially, the second constraint is omitted, and the remaining system (which is a generalised eigenvalue problem) is solved. Then, the full system is solved iteratively, the initial result being used as a starting value. Generally just a few (three or four) iterations are needed. For a geometric interpretation of the initial value, see ref. [36].

3.3. Multiple regions

In the second stage of direct segmentation we decompose any remaining smooth multiple regions. Many are defined by sweeping smooth profiles, which consist of tangential line-segments-and-circular-arcs or free-form profiles. Assuming translational or rotational symmetry has been detected in the given region, we can compute the best sweeping profile, as detailed in the next section. In other cases in step 5, we deal with *general* multiple regions. Fig. 4 shows a simple object, by means of which we illustrate in Section 5 how the topology of the final model is built. The object has not only a linear extrusion region which includes the five faces (TF1–TF5), but a general, multiple face region (SF1–SF7) as well. This latter has one toroidal, three planar, and three cylindrical faces, which were generated presumably in the original object by means of a blending operation, but due to the large radii used, these were not separated in the planarity filtering phase of reverse engineering.

The standard technique for segmenting regular surfaces is to analyse the signs of the estimated mean and Gaussian

curvatures [6]. The curvatures can be estimated based on the triangulation, or by fitting the best fit quadric at the given point. For noisy data, we apply a useful tool called *dimensionality filtering*. Taking the normal vectors of a multiple region mapped onto the Gaussian sphere, the normals of any planar region will be mapped into a point cluster; those of a linear extrusion or conical region onto a circular arc. Classifying the normal vector distributions on the Gaussian sphere as 0- or 1-dimensional, the planar and linear extrusion surfaces can be separated and detected [40].

The only problem left is when more than one surface of revolution exist within a smooth sub-region. In this case, we may identify all likely rotational axes using an approach suggested by Lukács [32]. The basic idea is that any four normal lines determine two axes which intersect all four lines, and which can be computed easily by a second-degree equation. By clustering the potential axes, we detect surfaces of revolution locally. This step concludes the direct segmentation procedure.

Returning to our segmented benchmark object in Fig. 3, the regions are shaded differently and detected simple surfaces have already been fitted. The multiple regions, the rotational axis of LEFT and the sweeping direction of the side face of RIGHT have also been computed and decomposed according to their smooth profile curves, i.e. we have generated a sequence of connected cylinders, cones, and tori for LEFT, and connected cylinders for RIGHT, respectively. The segmentation of the other simple test object can be seen in Fig. 5. These point regions are sufficiently stable to serve as an initial value for constrained fitting which enforces smoothness between adjacent surface elements. Our approach to constrained fitting is given in Ref. [4].

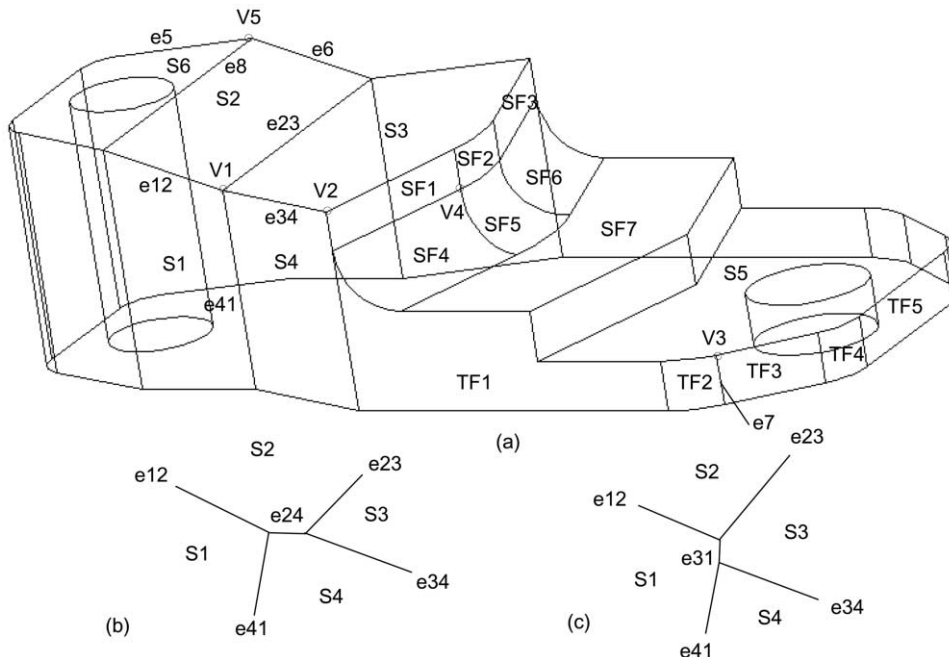


Fig. 4. Simple object.




Fig. 5. Segmenting the multiple smooth region of the simple object.

4. Constructing swept surface profiles

We continue with our approach to computing profile curves for swept regions. These are made up of smoothly connected straight line segments and circular arcs. The underlying smooth, multiple region has already been detected in the segmentation phase.

4.1. Guiding polygon

The *guiding polygon* is a key element in profile fitting. Approaches which are based only on point data may find it difficult to approximate the point set after projection. Due to the inaccuracies of the translational direction or rotational axis, we may obtain a point cloud which is ‘thick’, and thus it is hard to order the points, a step which is generally needed for approximating methods. An approach for thinning point clouds was reported recently [29]. We propose to work over the triangulation, which has the advantage that not only the point information, but their *connectivity* can be transformed into the plane of the profile.

Once we know the translational direction or the rotational axis, we can create a polygon which approximates the profile by intersecting the triangulation with a suitably chosen plane (perpendicular to the translational direction or containing the rotational axis as appropriate). In choosing a suitable plane, the problem is that many candidate profile planes do not contain the *entire* profile due to trimming of the swept face by adjacent faces—it may even be the case that *all* planes have this unfortunate property. This is why we have to apply a gradual approach to build up the complete polygon. We start at an arbitrary triangle of the swept region, chosen near the centre of the point set. We intersect it with the profile plane through its barycentre and extend the polygon by taking as many adjacent triangles intersected by that plane as possible. Of course, this may be just part of the guiding polygon and it may need to be further extended. After building the initial polygon, we check all triangles remaining in the region. Each triangle is

projected (or rotated) into the profile plane and its image is projected onto the polygon. If the current polygon does not fully cover the image, the polygon must be extended accordingly. In order to build up the polygon correctly, we do not traverse the triangles in an arbitrary order; we always take a triangle adjacent to a previously processed one.

If the distance between the two endpoints of the polygon becomes too small (say, smaller than the average distance of adjacent vertices), we close the polygon and stop the building process.

4.2. Constrained profiles

Of course, every profile can be approximated by a free-form curve, but many engineering applications require an explicit straight line–circular arc representation of the profile where applicable. While for free-form curves the guiding polygon provides a sensible parametrisation of the projected points, here it helps to segment the profile into circular arcs and straight line segments by applying a ‘recover and select’ procedure [31].

Circles are represented in Pratt’s form [35], which includes lines as special cases and not as singularities. After segmentation, circles whose radii are too great are replaced by straight lines. Then, the points are grouped with the help of the guiding polygon and either a circle or a straight line is fitted to each group, such that each is tangent to its neighbouring elements. There are two strategies to achieve this goal. One of them is to fit the elements sequentially, always satisfying the smoothness constraint with the previous element. This strategy may lead to a contradiction in the case of a closed profile; moreover, errors are accumulated: the first object to be fitted will be quite exact, but as we proceed we obtain worse and worse approximations. We prefer an alternative strategy which simultaneously approximates the points and satisfies the tangency constraints for all elements at once.

The equation of a circle written in Pratt’s form is the following:

$$A(x^2 + y^2) + Bx + Cy + D = 0,$$

where $B^2 + C^2 - 4AD = 1$. With the given normalisation condition, the expression on the left-hand-side is an approximation of the signed geometric distance of the point (x, y) from the circle, and consequently it is appropriate for a least squares minimisation. Taken in itself, this corresponds to a generalised eigenvector problem. As one can see, if $A = 0$, we obtain the equation of a straight line. The constraint expresses that the normal vector to the curve is a unit vector.

For constrained fitting, we need equations which express tangency between two circles or between a circle and a line. These are, for two circles:

$$(A_1 B_2 - A_2 B_1)^2 + (A_1 C_2 - A_2 C_1)^2 - (A_2 \pm A_1)^2 = 0;$$




Fig. 6. Reconstructed rotation profile.




Fig. 8. Circles of the rotation profile.

for a line and a circle:

$$B_1B_c + C_1C_c - 2D_1A_c \pm 1 = 0.$$

Based on these, a large system is set up, which minimises the square of the distances of the data points from the profile elements while at the same time satisfying a set of constraint equations. One may apply the general method of Lagrangian multipliers, or an alternative method of constrained fitting, as described in detail in a companion paper [4]. This solution requires an iterative process starting from initial values. These are computed during segmentation of the guiding polygon by independent unconstrained fits to each profile element.

Returning to the benchmark object, Fig. 6 shows the segmented profile curve of the LEFT component. In Fig. 7 the profile of RIGHT is shown. The construction elements, i.e. the full circles generating this profile, can be seen in Figs. 8 and 9. Having the profile curves, we can create the exact surface elements for each multiple face, which can directly be used in the final model building phase.

5. Building the boundary representation model

In order to build the B-rep model we have to compute all the edges and vertices explicitly. As discussed earlier, direct

segmentation has produced a set of disjoint regions. Each simple region has already been approximated by a single analytic surface. The smooth, multiple regions have been further segmented and approximated using the constrained fitting method, which has produced not only a set of analytic surfaces, but smooth internal curves as well, representing common boundaries between pairs of adjacent surface elements.

Although our two-step segmentation procedure looks complex, it was formulated to avoid the potential difficulties in computing edge curves. Surface–surface intersections work well to produce sharp edges whose existence was deduced during the first planarity filtering phase. Surface–surface intersections are likely to fail for smooth edges, where the adjacent surfaces are in principle tangential, due to inaccurate measured data and possible numerical inaccuracies of the surface fitting procedure. Nevertheless, smooth internal curves can be easily computed for linear extrusion surfaces and surfaces of revolution based on the vertices of the profile curve and the type of sweeping—extrusion or revolution. The internal curves within smooth, multiple regions are by-products of the constrained fitting procedure, and are




Fig. 7. Reconstructed translation profile.




Fig. 9. Circles of the translation profile.

explicitly available together with their points of intersection.

Let us return to our simple object shown in Fig. 4 to describe the procedure of creating B-rep models. Of course, this description cannot contain all the details, but it is sufficient to introduce the basic concepts and most important steps of the algorithm.

5.1. Fattening

Model building starts by creating a *region adjacency* data structure, which fundamentally reflects the topology of the final B-rep model. In the segmentation phase, triangles which were not unambiguously classified were discarded. Now, we reclassify these triangles and remove the gaps between the disjoint regions. This phase is called *fattening*, in which regions are extended outwards until all triangles have been classified. In one pass, those unclassified triangles are considered which have at least one common edge with the classified triangles. Take one of these triangles. If it has only one classified edge, the distance of the opposite vertex from the surface to which the edge belongs is computed. If it is within tolerance, the triangle is added to the region. If the triangle has more than one classified edge, the distances between the related surfaces and the opposite vertices are computed, and the closest region (assuming it is within tolerance) will acquire the triangle. If there are no further triangles which can be added because they satisfy the tolerance criterion, we must continue based only on the adjacency of the triangles. In this pass, all regions are simultaneously grown in a reasonably even manner.

By means of fattening we have practically created the topology of the object. Edges of the triangulation which are shared by two adjacent regions form polylines, from which we can deduce the existence of edges in the final model. Vertices of the triangulation which are shared by at least three regions and are located at the ends of the polylines correspond to vertices in the final model. Connected sets of triangles which logically all belong to the same region correspond to faces.

Special treatment is needed in the case of ‘very short’ edges and ‘very small’ faces, which presumably have arisen due to inaccurate determination of vertices with more than three edges. For example, take vertex $V1$ in Fig. 4. The common intersection point of the ideal surfaces $S1$, $S2$, $S3$ and $S4$ is quite likely to appear in the region adjacency structure in a different way. Due to numerical errors, fattening may detect $S2$ and $S4$, or $S1$ and $S3$ as neighbouring pairs, leading to the topologies shown in Fig. 4(b) and (c). However, the true topology must be consistent with the surface geometries, and this may differ from the configuration determined by fattening. This issue will be addressed later.

Another kind of special treatment is needed if the model is incomplete. Take, for example, edges $e5$ and $e6$ in Fig. 4, where the back faces adjacent to these edges are supposed to

be missing, or the partial reconstruction of an object based on a single view in Fig. 13. While typically, region boundaries are shared by two regions, in these cases there are no neighbouring triangles on the other side and *silhouette edges* must be generated to complete the model.

5.2. Edges and vertices

Computing edge curves is simple. In the case of a sharp edge, such as $e34$, we just call a surface–surface intersection routine to compute it, while if we have a smooth edge, such as $e7$, we just deduce the curve from the results of the constrained fitting phase, which guarantees that the smooth edge curve lies on both surfaces. For efficiency reasons, and to avoid ambiguities, we compute a bounding box for each edge, based on the related approximating polyline. We always leave processing of very short and silhouette edges to as late as possible as they are ill-defined.

Vertices where three surfaces meet are well determined. These can be determined by curve–surface intersection. For example, to compute $V2$, one option is to take $e34$, which is the common intersection curve of surfaces $S3$ and $S4$, and intersect it with surface $SF1$. If one of the edges is smooth, for example at vertex $V3$, special care is needed: here the smooth edge $e7$ shared by $TF2$ and $TF3$ must be intersected with the remaining third surface $S5$. Completely smooth vertices such as $V4$ are either determined during constrained fitting, or if necessary, by curve–curve intersection. Silhouette vertices, such as $V5$, are computed only approximately, constrained by the underlying edge curve $e8$ in our case.

If no constraints have been enforced during surface fitting, vertices with more than three edges need special care. Take the previous example of vertex $V1$. While the fattening procedure may tell us that the topology is as shown in Fig. 4(b), upon intersecting the various faces concerned, the geometry may produce the result shown in Fig. 4(c). We must be careful to take note of this and modify the topology accordingly. For example, moving towards the critical area on $e12$, one has to intersect edge $e12$ with $S3$ and $S4$; if the intersection point with $S4$ occurs ‘earlier’, we have configuration (b) and we insert a small edge $e24$, otherwise configuration (c) is correct and we have to insert edge $e13$. Of course, for complex edge configurations there are many more cases which may arise, but the principle of ordering the intersections helps to establish the correct topology. At a later stage, one of the basic tasks of beautification is to discover these short edges and explicitly enforce coincidence.

At the edge of the triangulation (if any), we must terminate the model by silhouette edges. In order to complete the edge loops of the corresponding faces, we must compute approximate edges by tracing the border polylines of the triangulation, for example, along edges $e5$ or $e6$. In the general case, we generate a piecewise linear approximation to the polyline and project it onto the underlying surface. Thus, a sequence of small edges lying on this face will be

obtained. If the face involved is a plane, an attempt is made to reconstruct the silhouette edge as a better composite curve made up of straight and circular segments, using the same methods as for profile reconstruction.

5.3. Stitching

We must now create the complete topological structure by stitching together the faces, edges and vertices. This is quite a straightforward process, since in the previous phases the consistency of the geometrical and topological entities has been assured. Taking an edge loop of a given face, the procedure applied guarantees that each real edge is shared by another edge of a neighbouring face and the related end vertices are identical. Thus, taking all edges of the loops of a given face, all adjacent faces can be stitched together. This operation is supported by most solid modelling kernels.

6. Blending surfaces

In the early planarity filtering phase we temporarily removed data points which belong to blending surfaces. Now, having already reconstructed the primary surfaces, we can reconstruct the blends. It is assumed that these were generated by constant-radius rolling-ball blending. We must determine the best approximation to the appropriate radius and just forward this information to the underlying solid modeller which will explicitly construct an appropriate blend and incorporate it into the B-rep model. Based on the region neighbourhood graph and the triangulation, we ignore point regions around vertices, and take the disjoint point sets which represent single edge blends between two given primary faces. Note that in this context the primary face can possibly be a smooth, multiple face; this makes no difference from the point of view of blending.

There are different approaches to estimating blend radii; these are thoroughly analysed in Ref. [27]. *Iterative spine* methods are based on the principle that the spine curve of a rolling ball blend is computed by intersecting offset primary surfaces. Having an initial radius estimate, the spine and the corresponding blend can be computed. After calculating the distances of the data points from this estimated blend, the radius can be corrected and a new spine computed; this is done iteratively until all the points are within tolerance. This method is particularly efficient in special cases when the intersection of the primary surfaces is a simple, analytic curve, but it also works for surfaces of arbitrary type.

Another approach is to compute a so-called *maximum ball* for each data point such that the ball contains the point and touches the two adjacent primary surfaces. Having the radius estimate for the maximum ball at each point, we compute an average for the whole blend. This method is also efficient and numerically stable, and it has the advantage



Fig. 10. Test object without small blends.




Fig. 11. Test object with small blends.

that it can be generalised to reproduce variable radius rolling ball blends.

We attach blend information to the edges, and let the modelling kernel incorporate the edge blends into the B-rep model. We do not reconstruct vertex blends in the current development phase, since in the case of small blends we only have a few data points. Vertex blends are considered to be by-products which automatically connect edge blends and are constructed by the solid modeller. Setback vertex blends provide a reasonable transition surface in the case of arbitrary edge blend configurations [39].

Fig. 10 shows the reconstructed benchmark object with sharp edges. The blended object can be seen in Fig. 11.

7. Implementation and experimental details

A prototype system for solid model reconstruction is under development at the Geometric Modelling Laboratory of the Computer and Automation Research Institute. A REPLICA laser scanner system (3D Scanners, London, UK) is used for data acquisition. The program is written in C++ and uses the VTK graphics library; it runs under LINUX and Windows NT. The ACIS solid modelling kernel [11] is used for solid model creation and blending.




Fig. 12. Segmentation of a single view point cloud.




Fig. 13. Solid object reconstructed from the single view.

The original simulated dataset for our benchmark test object shown in Fig. 1, contained 18,970 points. The size of the bounding box was $100 \times 50 \times 40$. The noise added to the sampled data points had standard deviation 0.05 units. In Fig. 1 the point set was decimated to 8238 points with 16,432 triangles. The estimated rotational axes of left and the translational direction of right were $(0.0028, 0.9999, -0.0008)$ and $(0.0232, 0.0052, 0.9997)$; compare these with the exact vectors of $(0, 1, 0)$ and $(0, 0, 1)$. The radius of the large cylinder of LEFT was computed as 30.1478 (cf. 30); that of RIGHT was computed as 26.876 (cf. 27). The radii of the profile elements were reconstructed with reasonable accuracy. For example, for the rotational profile, the computed radii are 4.02 (cf. 4), 3.84 (cf. 4), 4.31 (cf. 4), 3.97 (cf. 4). The blend reconstruction also gave reasonable values, for example, for the top faces of MIDDLE and RIGHT we obtained radii of 2.03 (cf. 2) and 2.01 (cf. 2); along their intersection the radius was 1.01 (cf. 1).

To illustrate the difficulties of dealing with real measured data, another object was selected; the segmentation of the corresponding data is shown in Fig. 12. This is a single view dataset, and so the stitched model is incomplete. Compare




Fig. 14. Another view of the reconstructed solid object.




Fig. 15. Solid object reconstructed from multiple views.

figures viewed from the scanning direction (Fig. 13) and another slightly rotated view (Fig. 14). The latter may look horrible, but it is a valid ACIS model. Of course, if we start with a suitable multiple view point cloud representing the complete solid, the artefacts disappear, as shown in Fig. 15.

8. Conclusion

A methodology for reverse engineering conventional solid objects has been presented, producing boundary representation solid models, which can be directly used in mechanical engineering CAD/CAM. The complex nature of the procedure was illustrated by describing the consecutive algorithmic steps and their links. The importance of reproducing the regularity of the objects—simple surfaces, translational and rotational symmetries, constraints—cannot be overemphasised. The current prototype implementation shows that for reasonably dense and accurate measured data, solid reconstruction is possible with minimal user interaction by setting only certain threshold and tolerance values which control the various steps of the algorithm. Current and future efforts will be mainly directed towards reliably segmenting complex multiple regions,

automatically deriving constraints, constrained fitting and model beautification.

Acknowledgements

The solid reconstruction project started in 1995 within the framework of an EU supported COPERNICUS grant (RECCAD 1068). The authors highly appreciate the contribution of the GML research team, including Géza Kós, Gábor Lukács, Péter Salvi and László Andor from CADMUS Ltd. This research was also supported by the National Science Foundation of the Hungarian Academy of Sciences (OTKA, No. 26203) and by the UK EPSRC (Grant GR/M 78267).

References

- [1] Alboul L, van Damme R. Polyhedral metrics in surface reconstruction: tight triangulations. In: Goodman T, Martin RR, editors. *The mathematics of surfaces VII*, Information Geometers Ltd Winchester, 1997. p. 309–36.
- [2] Bajaj C, Bernardini F, Chen J, Schikore D. Automatic reconstruction of 3D CAD models. In: *Proc Int Conf on Theory and Practice of Geometric Modeling II*, Blauberger, Germany, October 1996.
- [3] Benkő P, Lukács G, Martin RR, Várady T. Algorithms to create B-rep models from measured point clouds. *Geometric Modelling Studies*, GML 1998;6:1998.
- [4] Benkő P, Andor L, Kós G, Martin RR, Várady T. Constrained fitting in reverse engineering. *Computer Aided Geometric Design 2000* (submitted).
- [5] Bergevin R, Soucy M, Gagon H, Laurendeau D. Towards a general multi-view registration technique. *IEEE PAMI* 1996;18(5):540–7.
- [6] Besl PJ. *Surfaces in range image understanding*. Springer-Verlag, 1988.
- [7] Besl PJ, Jain RC. Segmentation through variable-order surface fitting. *IEEE PAMI* 1988;10(2):167–92.
- [8] Bolle RM, Vemuri BC. On three-dimensional surface reconstruction models. *IEEE PAMI* 1991;13(1):1–13.
- [9] Chen X, Hoffman CM. On editability of feature-based design. *Computer Aided Design* 1995;27(12):905–14.
- [10] Chivate PN, Jablowski AG. Solid-model generation from measured point data. *Computer Aided Design* 1992;25:587–600.
- [11] Corney J. *3D modeling using the ACIS Kernel and Toolkit*. Chichester: John Wiley & Sons, 1997.
- [12] Dorai C, Wang G, Jain AK, Mercer C. Registration and integration of multiple object views for 3D model construction. *IEEE PAMI* 1998;20(1):83–89.
- [13] Eck M, Hoppe H. Automatic reconstruction of B-spline surfaces of arbitrary topological type. *Computer Graphics (SIGGRAPH 96)*, vol. 30, 1996.
- [14] Edelsbrunner H, Mücke E. Three-dimensional alpha shapes. *ACM TOG* 1994;13(1):43–72.
- [15] Faugeras OD, Hebert M. The representation, recognition and locating of 3D objects. *Int J Robotics Research* 1986;5(3):27–52.
- [16] Fitzgibbon AW, Eggert DW, Fisher RB. High-level CAD. model acquisition from range images. *Computer Aided Design* 1997;29(4):321–30.
- [17] Greiner G, Hormann K. Interpolating and approximating scattered 3D data with hierarchical tensor product B-splines. In: Le Méhauté A, Rabut C, Schumaker LL, editors. *Surface fitting and multiresolution methods*, Vanderbilt University Press Nashville & London, 1997. p. 163–72.
- [18] Guo B, Menon J, Willette B. Surface reconstruction using alpha shapes. *Computer Graphics Forum* 1997;16(4):177–90.
- [19] Hebert M, Ponce J. A new method for segmenting 3D scenes into primitives. In: *Proc Sixth Int Conf on Pattern Recognition*, München, 1982. p. 836–8.
- [20] Heckbert PS, Garland M. Survey of polygonal surface simplification algorithms. *Course Notes, Course. 25, SIGGRAPH 97*, 1997.
- [21] Hoover A, Goldgof D, Bowyer KW. Extracting a valid boundary representation from a segmented range image. *IEEE PAMI* 1995;17(9):920–4.
- [22] Hoover A, Jean-Baptiste G, Jiang X, Flynn PJ, Bunke H, Goldgof DB, Bowyer K, Eggert DW, Fitzgibbon A, Fisher RB. An experimental comparison of range image segmentation algorithms. *IEEE PAMI* 1996;18(7):673–89.
- [23] Hoppe H, DeRose T, Duchamp T, McDonald J, Stuetzle W. Surface reconstruction from unorganised points. *Computer Graphics (SIGGRAPH 92)*, vol. 26, 1992. p. 71–8.
- [24] Hoschek J, Dankwort W, editors. *Reverse engineering*. Stuttgart: Teubner, 1996.
- [25] Hoschek J, Kaklis P, editors. *Advanced course on FAIRSHAPE*. Stuttgart: Teubner, 1996.
- [26] Hoschek J, Dietz U, Wilke W. A geometric concept of reverse engineering of shape: approximation and feature lines. In: Dæhlen M, Lyche T, Schumaker LL, editors. *Mathematical methods for curves and surfaces II*, Vanderbilt University Press Nashville & London, 1998. p. 253–62.
- [27] Kós G, Martin RR, Várady T. Recovery of blend surfaces in reverse engineering. *Computer Aided Geometric Design* 2000;17:127–60.
- [28] Kós G. An algorithm to triangulate surfaces in 3D using unorganised point clouds. *Geometric Modelling Studies*, GML 1999/3; Computer and Automation Research Institute, Budapest, 1999 (to appear in *Computing* 2000).
- [29] Lee IK. Curve reconstruction from unorganised points. *Technical Report No. 55*, Institut für Geometrie, Technische Universität Wien, 1998.
- [30] Krsek P., Pajdla T., Hlaváč V., Martin RR. Range image registration driven by a hierarchy of surface differential features. In: *22nd Workshop of the Austrian Association for Pattern Recognition*, Eds. M Gengler, M Prinz, F Schuster, 175–183, Osterreichische Computer Gesellschaft, 1998.
- [31] Leonardis A, Jaklic A, Solina F. Superquadrics for segmenting and modeling range data. *IEEE PAMI* 1997;19(11):1289–95.
- [32] Lukács G, Marshall AD, Martin RR. Faithful least-squares fitting of spheres, cylinders, cones and tori for reliable segmentation. In: Burkhardt H, Neumann B, editors. *Computer Vision-ECCV 98*, Lecture notes in computer science, Springer-Verlag, 1998.
- [33] Marshall AD, Lukács G, Martin RR. Robust segmentation of primitives from range data in the presence of geometric degeneracy. *IEEE PAMI* 2000;23(3):304–314.
- [34] Parvin B, Medioni G. B-rep, from unregistered multiple range images. In: *Proc Int Conf Robotics and Automation*, Nice, 1992. p. 1602–7.
- [35] Pratt V. Direct least-squares fitting of algebraic surfaces. *Computer Graphics (SIGGRAPH 87)*, vol. 21(4), 1987. p. 145–52.
- [36] Pottmann H, Randrup T. Rotational and helical surface approximation for reverse engineering. *Computing* 1998;307–22.
- [37] Sapidis NS, Besl PJ. Direct construction of polynomial surfaces from dense range images through region growing. *ACM TOG* 1995;14(2):171–200.
- [38] Várady T, Martin RR, Cox J. Reverse engineering of geometric models—an introduction. *Computer Aided Design* 1997;29(4):255–68.
- [39] Várady T, Hoffmann CM. Vertex blending: problems and solutions. In: Dæhlen M, Lyche T, Schumaker LL, editors. *Mathematical methods for curves and surfaces II*, Vanderbilt University Press Nashville & London, 1998. p. 501–28.

- [40] Várady T, Benkő P, Kós G. Reverse engineering regular objects: simple segmentation and surface fitting procedures. *Int J of Shape Modeling* 1998;4:127–41.
- [41] Werghi N, Fisher R, Robertson C, Ashbrook A. Object reconstruction by incorporating geometric constraints in reverse engineering. *Computer Aided Design* 1999;31:363–99.
- [42] Langbein FC, Mills BI, Marshall AD, Martin RR. Finding approximate shape regularities in reverse engineered solid models bounded by simple surfaces. To appear, *Proc. ACM Solid Modelling* 2001.
- [43] Mills BI, Langbein FC, Marshall AD, Martin RR. Approximate Symmetry Detection for Reverse Engineering. To appear, *Proc. ACM Solid Modelling*, 2001.