An Improved Algorithm to Extract Surfaces from Complete Range Descriptions

C. Robertson, R. B. Fisher, N. Werghi, A. P. Ashbrook
Vision Group, Institute for Action, Perception and Behaviour,
Division of Informatics, University of Edinburgh,
Edinburgh, UK EH1 2QL
craigr@dai.ed.ac.uk

Keywords: range data segmentation, quadric surface extraction

Abstract

Reverse-engineering a machined part to generate a CAD model requires range data to be collected and registered from many views then segmented into surface primitives. Most previous research on surface extraction has concentrated on extracting surfaces from only one view at a time and then registering them afterwards. This leads to alignment problems from combining partial surface fragments in order to produce complete models. We have avoided these problems by implementing a 3D segmentation algorithm to extract surfaces directly from fully registered range data clouds. This paper reports on improvements to the algorithm which have had the effect of reducing errors in the surface parameter estimation as well as increasing the overall robustness of the technique.

1 Introduction

Reverse-engineering a machined part to generate a CAD model requires range data to be collected and registered from many views then segmented into surface primitives. There are two main approaches to doing this: the first is to extract features (surfaces) from individual range data sets and then register these features in several views; the second is to register the 'raw' range datasets and then extract surfaces from the resultant data-cloud.

There are several problems associated with combining surfaces from single orthogonal views. The first is how to accurately register the occuring features into a single reference frame. It is possible to align the views provided three independent non-parallel planes can be found in each pair of views. Two planes provide orientation alignment but allow one translational degree of freedom, which is resolved by the third plane. Faugeras [5] gives standard algorithms for transformation estimation but we have found that this provides only limited accuracy. Our models are therefore typically accurate only to about 1-2 mm

when using a range sensor that has 0.15 mm measurement standard deviation. To improve registration variations of the iterated closest point algorithm [1, 23] are often used to provide fine alignment using the raw range data after the feature-based alignment provides an initial coarse registration.

The problem of registering multiple range datasets has been well researched, for example by [1, 2], in the past few years. The result of this research is that it is now possible to obtain quite good registration of clean range image datasets of smooth curved and developable surfaced objects. However, even if the views can be accurately aligned there is still the problem of stitching together the features. This is a rather complex problem because individual views might only provide a fragment of a surface and even a pair of views might not provide a complete view. One might use a symbolic surface-stitching algorithm to produce new surfaces, for example Orr [11] investigated methods for producing a combined planar surface from multiple overlapping views. His algorithm required complex reasoning in order to construct the boundary of a combined surface from the boundaries of the individual surfaces, particularly because positioning errors require one to use statistical tests to determine if vertices observed in the two views are the same. The problems that have been encountered derive mainly from errors in registration: it is difficult to determine when features found on surface boundaries arise from segmentation limitations, statistical variations or nearby real features. Orr overcame some of the problems using a Mahalanobis distance test to assess the likelihood of statistical variations, but problems remained with the reliability of estimating surface statistics and the resolution of the remaining alternative merging hypotheses.

Work has been done previously on extracting geometric surfaces from full 3D datasets, of which [6] is a good example, although it addressed mainly planar surfaces. Successful previous research on extracting descriptions from segmented range images has included [13, 14, 15] who explored combining both surface patches (refitting surfaces to regions that overlap in registered multiple views) and combining polygonal surfaces. An alternative is to fit completely covering spline meshes to a polygonalized 3D dataset, once the views have been registered and combined [10, 9].

The research presented here avoids these problems by adapting a single view quadric surface range segmentation program to extract surfaces and thus complete object models directly from fully registered 3D range datasets. It can be applied to both industrial parts with simple developable surfaces and more complex surfaces.

The research presented here builds on the approach concept-tested by Fisher *et al* [20]. It approach improves on this work by handling curved surfaces better, applying robust methods for surface growing and leaving the possibility for later surface fit optimization using *a priori* geometric constraints by both emergent algorithms [18] and variations on the Levenberg-Marquardt algorithm [19]. Results for parts which have been analysed have shown a reduction in the errors by 50% and together with the further optimization techniques, it seems likely that even this can be improved.

2 Surface Extraction

The method for complete surface extraction consists of the following distinct stages:

1. Multiple view surface registration.
2. Surface polygonalization.
3. Local curvature estimation and shape classification.
4. Surface growing.
5. Assignment of range image vertices to surfaces.
6. Final surface fitting and specialized fitting.

2.1 Surface Registration

The algorithm used for multiple view surface registration is described in [1]. The algorithm registers multiple range images using a modified *iterated closest point* algorithm [1, 4] and produces a cloud of range (x, y, z) values registered in a common reference frame. The algorithm assumes that every point has been observed in at least two range images, and thus avoids the problem of identifying points without a correspondence, meaning that no distance threshold is needed.

2.2 Surface Polygonalization

With range data acquired from a single viewpoint there is usually a local topology in the form of a regular grid imposed by the mechanical structure of the range sensor. With multiple registered range images the range data usually forms a cloud of points about the true surface. This means that getting an initial approximation of the surface shape is difficult and it is hard to determine point adjacency on the surface during the surface growing phase. Our solution to this problem was to tessellate the surface, giving a regular 3D mesh, then to perform surface shape extraction and growing using the centroids of the mesh polygons. The initial surface polygonalization is performed using the Hoppe *et al* algorithm [8]. The purpose of this polygonalization is to construct local topology between the range points which may then be used as part of the curvature classification and surface growing processes. Shown in figure 1(a) is a typical polygonization, of a half cylinder object. The regularity of the mesh can be seen clearly. Figure 1(b) shows the rendered polygon mesh for a partly segmented sloping faced cube object with the polygon patch normals rendered at the polygon centroids. The instability of the normals nearing the fold edges can also be seen from this figure. It is important that the polygon representation preserves the statistical shape characteristics of the original data points as well as possible.

2.3 Local Surface Curvature Estimation

Given the local topology, it is then easy to estimate the local surface curvatures at each point with a local surface fitting algorithm. This provides an

estimate of the principal curvatures which could be used in a (Mean, Gaussian) curvature classification process. However, this process would not work well because surface triangulation greatly reduces the number of data points (that is vertices) available for calculation when compared to a density of information available from single view. Taubin [16] and Chen and Schmitt [3] give linear and quadratic complexity algorithms (respectively) for estimating the local principal curvatures at the vertices of a polyhedral mesh. Unfortunately, polygonalized surfaces formed from merged datasets are locally rather noisy so the curvature estimates are not particularly accurate. Our experience has also shown that the surface fitting algorithm has a more significant impact than the initial curvature classification. So, because our surface fitting algorithm uses a general quadric form rather than different forms specialized for the different shape classes, we only classify polygons into three classes, {planar, curved, edge}.

Shape classification estimates the local curvedness at each polygon by finding the maximum angle θ_{max} that nearby connected surface polygons turn away from the current polygon. This angle is found by examining the angle between polygon normals in a neighbourhood of distance N polygons (default is $N = 2$) about the current polygon. The maximum angle gives an indication of how curved the surface is and forms the basis for the initial labeling of polygons into different surface shape classes:

If $\theta_{max} < \tau_{plane}$ (5.0 deg)	Plane
If $\theta_{max} > \tau_{edge}$ (10 deg)	Edge
Else	Curved

These parameters were selected by hand-tuning, but were neither very stable (always reliable classifications) nor unstable (quite varying classifications). Changing the parameters gradually increases the numbers of misclassified polygons, but this did not cause problems for the surface fitting algorithm. Using a surface growing algorithm, it is better to over-segment but still produce patches that are part of a single underlying surface which can then be grown together. Extreme over-segmentation should be avoided otherwise the initial surface patches are too small to get a decent fit. A finer mesh can help resolve the problem at the cost of increased computation. On the other hand, under-segmentation will produce initial regions combining portions of several surfaces whose initial fit will be bad, and will give parameters that do not lead to viable surfaces during surface growing.

Since this approach is viewpoint independent, depth discontinuities do not exist, so only surface orientation (fold edge) discontinuities are detected. Topologically connected polygons with the same shape class are joined to form extended surfaces which are the seed surfaces for surface growing.

2.4 Surface Growing

Complete surfaces are found by an iterative surface growing process that starts from the initial seed regions formed by grouping polygons which have the same

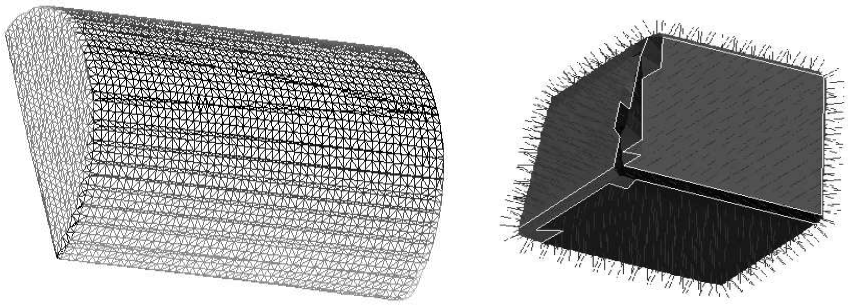

(a) Example Polygonization

(b) Example Polygonization showing Surface Normals

Figure 1: Surface Meshes

curvature type. The model for the surfaces describing these regions is the general quadric form. To establish the surface limits, the algorithm adds and deletes members from lists of polygons and uses the centroids of these polygons for fitting. Shown in fig.2 are two regions. In this iteration, region 1 may expand into neighbouring polygons from region 2 which are directly neighbouring it provided a set of criteria are fulfilled (distance to surface and estimated normals similar), depending upon the run-time shape classification for region 1. In the

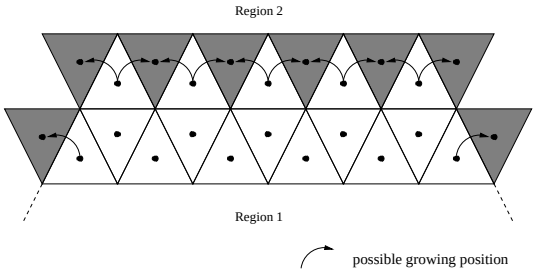

Figure 2: Local surface growing

final fitting pass, the cloud of raw range points are assigned to the surfaces and the surface is then estimated. The key ideas behind the algorithm are:

- There is competition between surfaces for polygons.
- The algorithm works iteratively between extending the surface shape and refitting the surface model to the grouped polygons.
- Small surfaces are deleted and their polygons are then available for combining into other surfaces.
- Only planes are extracted on the first pass, as their curvatures are less stable.

- After a surface is fitted, polygons that do not lie within $\sigma_{noise}\sqrt{2}$ of the fitted surface are returned to the pool of unfitted polygons. σ_{noise} is the estimated standard deviation of the polygon patch centroid from the true surface.
- All of the raw range points from the registered range data are assigned to surfaces and a final surface fit is then performed. If required, specialized degenerate quadric surface fitting is performed using constrained optimization methods outlined in [18].

A pseudo-code description of the reconstruction algorithm follows.

At each of the N passes there is a set of surfaces R (initially created by connected component analysis on the surface) that are being grown by the algorithm. Each surface consists of a set of edge connected mesh polygons and its implicit surface parameters. There is also a pool of mesh polygons that are not yet associated with any surface.

```
For each of N passes
  For each current surface R
    Do {
      If current surface too small ( < 50 polygons
      for example),
        then move polygons into pool
      If first pass then only extract planes
      Fit quadric surface to R
      Fit plane surface to R
      If surface fit fails, then move polygons into pool
      If second or later pass, then
        Move polygons from R into the pool that are
          further than σ_noise √2 from R's surface.
        If R is classified as a plane
          Add the polygons from the pool that :
          1. are adjacent to R
          2. have surface normal within ψ degrees
          (eg. 8 degrees) of the estimated surface
          normal
          3. closer than σ_noise √2 of
          R's surface.
        If R is classified as a quadric:
          Add the polygons from the pool that are :
          1. adjacent to R
          2. closer than σ_noise √2 of R's surface.
        Let B be the number of added polygons
      } while (B > 0)
```

Not all polygons in the mesh will necessarily be added to one of the regions. This is due to the two robust operations where polygons are removed if they are not sufficiently close to the surface fit.

2.5 Surface Model Selection and Parameter Estimation

Central to the surface fitting process is a least square shape parameter estimation process that determines both the surface type and the initial surface shape parameters. The key ideas of the shape parameter estimation are:

- A least-square error criterion is used.
- The fitting is for general quadrics (10 independent parameters: $ax^2 + by^2 + cz^2 + dxy + exz + fyz + gx + hy + iz + j = 0$). The parameters are estimated initially using Taubin's method [17].
- The plane fitting uses a standard least square method. If $(x_i, y_i, z_i)^T$ is the i^{th} data point, then let $\vec{p_i} = (x_i, y_i, z_i, 1)^T$. Then the parameter vector $(a, b, c, d)^T$ of the plane model $ax + by + cz + d = 0$ is the eigenvector with smallest eigenvalue from the scatter matrix $\sum_i \vec{p_i}\vec{p_i}^T$.
- The selection between the different model shape classes is based on minimizing the surface fit error, and the class j with the smallest e_j is chosen.

  $$e_{surface} = \Sigma\delta_j^2$$

  δ_i is the geometric error distance between the model and the polygon center-of-mass. In the case of the quadric, this is approximated using the approximation suggested by Taubin [17].

The pseudo-code for the surface fitting algorithm follows:

Get e_{plane} from plane fit
Get $e_{quadric}$ from quadric fit
If plane and quadric fit successful (ie. least-square
scatter matrices are full rank)
  Choose surface type with best fit
Else if plane fit successful then choose Plane type
Else if quadric fit successful then choose
Quadric type
  Try specialised degenerate surface fitting.
  (We currently employ cylindrical and spherical
  surface fitters.)
Else return NoFit

2.6 Range Data Point Assignment and Final Fitting

In the final stage points from the original range data are assigned to surfaces based on their closest geometric distance to all the available surfaces.

The pseudocode for the point assignment is as follows:

For all M points
  Find geometric distance to all fitted surfaces
  Let closest distance be d_{least} and second
  closest be d_{second}
  If abs($d_{least} - d_{second}$) <0.5mm then reject
  the point

else
add the point to the scatter matrix of the
best surface

After these assignments, the final fitting is performed for all the surfaces. In the results shown in section 3, a value of 0.5mm was found empirically. This value is to minimize the number of mis-assigned points due to intersections of what are essentially infinite theoretical surfaces. If any ambiguity exists, the point is rejected. This also reduces the possibility that points which lie at the intersection of planes (fold edges) are added to the incorrect plane.

3 Results

3.1 Typical Segmentations

Here we show the results of the segmentation process on three machined parts. In figure 3 we see a view of the combined range data for parts with planar, cylindrical and conical surfaces (UFO, SHOE and BED). The boundary is a little ragged because the Hoppe algorithm [8] was used to tessellate the surface and is sensitive to local surface noise, resulting in small outlying polygonal patches. Figure 4 show views of the segmented surfaces with the patch boundaries clearly visible. It should be noted that the polygons which lie between boundaries are the result of the robust application of the surface growing. They have not been added to the surfaces because either their distance from the estimated surface was too large or they have local normals that lie outside the specified range. One can see a little raggedness where the cylindrical surface has the smooth join with the planar surface and improving this is a topic of future research. Otherwise, there is no observable detail to indicate any flaws with the method, the segmentation algorithm has worked well on these and several other parts. Typical processing times, for example the BED part, are as follows: initial loading, grouping, fitting and surface growing (for 5477 polygons) : 34 seconds; final point assignment (for 40107 points) : 90 seconds. Some numerical results are given below.

3.2 Statistics

Based on the geometric features extracted from several objects, some accuracy statistics are (N is the number of instances of this relationship encountered in the experiments):

- mean angular error between planes is 0.12° (min = 0.01°, max = 0.38°, $N = 24$).
- mean cylindrical radius error is 0.20 mm (min = 0.02 mm, max = 0.33 mm, $N = 3$).
- mean parallel plane separation is 0.24 mm (min = 0.03 mm, max = 0.46 mm, $N = 6$).

- mean angle error between cylinder axes and an orthogonal plane is 0.38° (min = 0.00°, max = 1.2°, $N = 5$).

It should be noted that our range sensor produces range data with standard deviation about $\sigma = 0.15$ mm nearly uniformly over the whole field of view. The geometric accuracy is also greatly affected by the success of the surface registration algorithm [1]. Hence, these accuracy figures are really a function of both processes. The application of the new methods and revised process give results which improve the errors on the surface fitting by around 50% when compared to the previous algorithm [20].

4 Conclusions and Further Work

This paper has presented a segmentation algorithm that has the advantage of avoiding the problem of fusing symbolic surface descriptions when producing a description of all sides of a part in a common reference frame. It depends primarily on having registered 3D dataset to provide the accurate relative positioning of the surface points, from which our quadric surface extraction algorithm can then accurately extract surfaces.

Work is currently underway to improve the overall robustness by adding extra routines to the method:

- Addition of further robust quadric growing heuristics.
- Assessment of robustness using many subsampled meshes.
- Integration of optimization based on *a priori* constraints [19, 18].
- In future work a more elegant solution to point assignement will be adopted, with the introduction of point normals for range data points. This will improve assignment for final fitting since the point normal can be compared to the theoretical surface normals at that point. This will greatly reduce ambiguity.

Since it seems that the polygonalization algorithm is particularly important in the first phase (although less so in the final fitting phase) we are currently looking into methods for mesh reduction by decimation, to reduce the workload in this phase, as well as better surface approximation techniques.

Acknowledgements

The work presented in this paper was funded by a UK EPSRC grant GR/H86905.

References

[1] D. Eggert, A. W. Fitzgibbon, R. B. Fisher. “Simultaneous registration of multiple range views for use in reverse engineering”, Proc. Int. Conf. on Pat. Recog., pp 243–247, Vienna, Aug. 1996.

[2] H. Gagnon, M. Soucy, R. Bergevin, and D. Laurendeau. "Registration of multiple range views for automatic 3-D model building", Proc. IEEE Comp. Soc. Conf. on Computer Vision and Pattern Recognition, Seattle, WA, pp 581-586, June 1994.

[3] X. Chen and F. Schmitt. "Intrinsic surface properties from surface triangulation", Proc. 2nd European Conf. on Computer Vision, (ed) G. Sandini, pp 739–743 St. Margherita Ligure, Italy, May 1992.

[4] D. W. Eggert, A. Lorusso, R. B. Fisher, "Estimating 3-D rigid body transformations: a comparison of four major algorithms", Mach. Vis. and Applic., Vol 9, pp 272–290, 1997.

[5] O. D. Faugeras. Three-dimensional computer vision : a geometric viewpoint, Cambridge, Mass. MIT Press, 1993.

[6] O. D. Faugeras, M. Hebert. "A 3-D Recognition and Positioning Algorithm Using Geometric Matching Between Primitive Surfaces", Proceedings 8th Int. Joint Conf. on Artificial Intelligence, pp 996-1002, 1983.

[7] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. Goldgof, K. Bowyer, D. Eggert, A. Fitzgibbon, R. Fisher. "An Experimental Comparison of Range Segmentation Algorithms", IEEE Trans. Pat. Anal. and Mach. Intel., 18(7), pp 673–689, July 1996.

[8] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stuetzle. "Surface Reconstruction from Unorganized Points", Computer Graphics, 26(2), pp 71–78, 1992.

[9] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction, Proc. SIGGRAPH, pp 295-302, 1994.

[10] V. Krishnamurthy and M. Levoy. Fitting smooth surfaces to dense polygon meshes. Proc. SIGGRAPH: computer vision, pp 313-324, 1996.

[11] M. J. L. Orr, J. Hallam, R. B. Fisher. "Fusion through Interpretation", Proc. 2nd European Conf. on Computer Vision, (ed) G. Sandini, pp 801–805, St. Margherita Ligure, Italy, 1992.

[12] M. Oshima, Y. Shirai. "Object Recognition Using Three-Dimensional Information", Proc. 7th Int. Joint Conf. on Artificial Intelligence, pp 601-606, 1981.

[13] V. Sequeira, J.G.M Gonçalves, M. I. Ribeiro. 3D Environment Modelling Using Laser Range Sensing. J. of Robotics and Autonomous Systems, 16(1), pp 114-127, 1995.

[14] V. Sequeira, J. Gonçalves, M. I. Ribeiro. 3D Reconstruction of Indoor Environments. Proc. 3rd Int. Conf. Image Proc., Lausanne, Switzerland, pp. 17P2, September 16-19, 1996.

[15] V. Sequeira. Active Range Sensing for Three-Dimensional Environment Reconstruction. Eur. Commission Joint Research Centre Tech Note I.97.26, Feb 1997.

[16] G. Taubin. "Estimating the tensor of curvature of a surface from a polyhedral approximation", Proc. 5th Int. Conf. on Computer Vision, pp 902-907, 1995.

[17] G. Taubin. "Estimation of planar curves, surfaces and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation", Proc. IEEE PAMI, 13(11), pp1115–1138, 1991.

[18] C. Robertson, R.B. Fisher, D. Corne, N. Werghi, A.P. Ashbrook, "Investigating Evolutionary Optimisation of Constrained Functions to Capture Descriptions from Range Data", Advances in Soft Computing - Engineering Design and Manufacturing. R. Roy, T. Furuhashi and P. K. Chawdhry (Eds), Springer, 1999, pp 455–466.

[19] N. Werghi, R. B. Fisher, A. Ashbrook, C. Robertson, "Improving model shape acquisition by incorporating geometric constraints", Proc. British Machine Vision Conference BMVC97, Essex, pp 520–529 September 1997.

[20] R. B. Fisher, A. W. Fitzgibbon, D. Eggert, "Extracting Surface Patches from Complete Range Descriptions", Proc. Int. Conf. on Recent Advances in 3-D Digital Imaging and Modeling, Ottawa, Canada, pp 148-155, May 1997.

[21] G. Lukacs, R. Martin, D. Marshall,"Faithful Least-Squares Fitting of Spheres, Cylinders and Tori for Reliable Segmentation", Computer Vision, Proceedings of the 5th European Conference on Computer Vision, Feiburg, Germany, June 1998, Springer Verlag, Berlin, 1998.

[22] R. Benjamaa, F. Schmitt,"A Solution for the Registration of Multiple 3D Point Sets Using Unit Quaternions", Computer Vision, Proceedings of the 5th European Conference on Computer Vision, Feiburg, Germany, June 1998, Springer Verlag, Berlin, 1998.

[23] A. E. Johnson, M. Hebert,"Surface Registration by Matching Oriented Points", Proceedings of the International Conference on Recent Advances in 3-D Digital Imaging and Modeling 1997, Ottawa, Canada, IEEE Computer Society Press, Los Alamitos, California, USA, 1997.

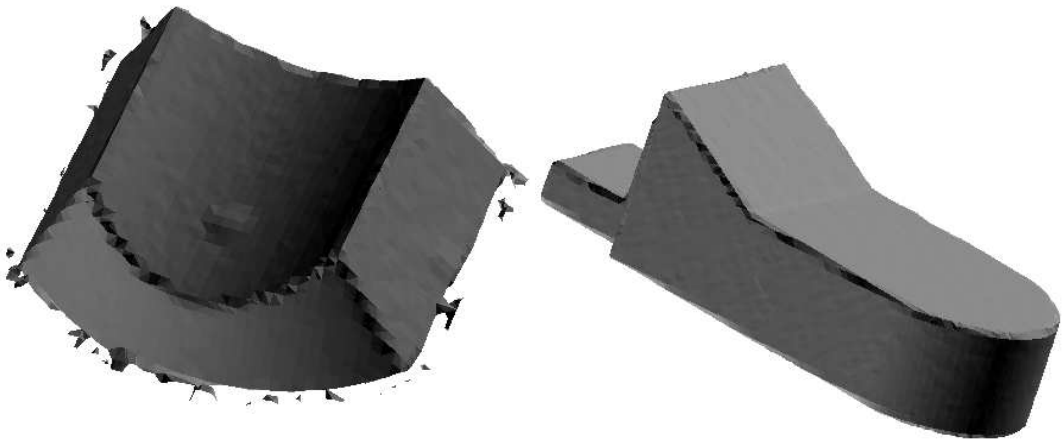

UFO Mesh

SHOE Mesh

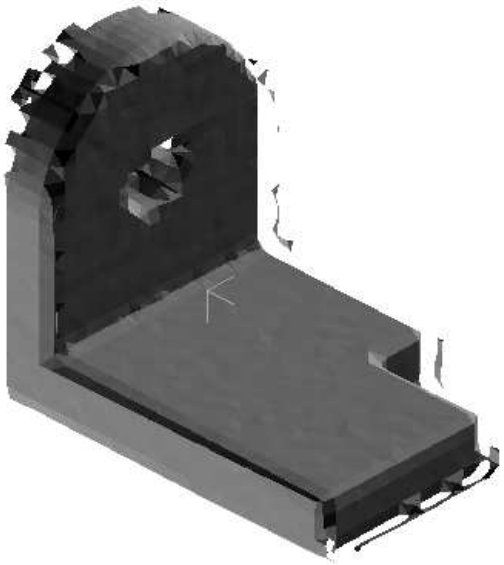

BED Mesh

Figure 3: Example Meshes - Notice the 'feathering' effect produced by the meshing algorithm. This is partially due to outliers in the data. The meshes are cosine shaded relative to a set of light sources.

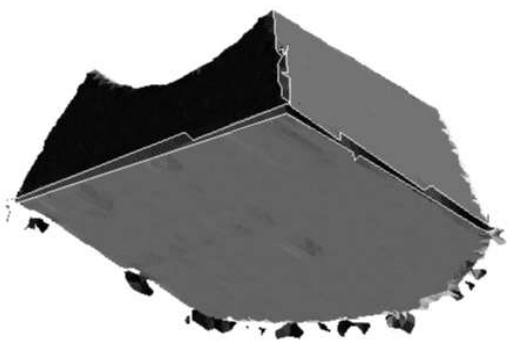

UFO

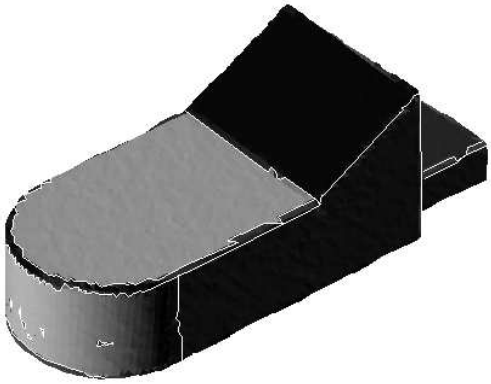

SHOE

BED

Figure 4: Segmented Meshes - The polygon meshes are coloured by their surface number and then cosine shaded as before. Extreme edges of the regions are shown by the white lines.