Applying knowledge to reverse engineering problems

Robert B. Fisher
Division of Informatics, University of Edinburgh
rbf@dai.ed.ac.uk

Abstract

This paper summarizes recent research at Edinburgh University on applying domain knowledge of standard shapes and relationships to solve or improve reverse engineering problems. The problems considered are how to enforce known relationships when data fitting, how to extract features even in noisy data, how to get better shape parameter estimates and how to infer data about unseen features.

1. Introduction

Traditional processes for reverse engineering objects and structures from 3D datasets have been initially data (*e.g.* triangulated models) and parametric surface (*e.g.* quadric surface) driven. These approaches has been successful for simple parts, but have resulted in reconstructions that have 'frozen-in' errors. Typical errors are surfaces not at the correct relative positions or artifacts arising from noisy or missing data.

For several years our research group has been exploring 'knowledge-based' techniques to help overcome these and other problems. The underlying theme behind this set of techniques is the exploitation of general knowledge about the domain of objects being reconstructed. The reconstruction process is not "model-based" reverse engineering, as then there would be no point to building the models - this would not be "reverse engineering". On the other hand, the knowledge is not arbitrary, because the objects that humans construct are not arbitrary: the shapes of most normal objects follow standard conventions arising from tradition or utility.

We argue that exploiting this extra knowledge allows improved reverse engineering. This paper presents several different examples of the general approach. Without apology we cite only our publications as pointers to more complete presentations of the material summarized here. Of course, the full publications contain proper citations and can be found at: http://www.dai.ed.ac.uk/homes/rbf/publications.html.

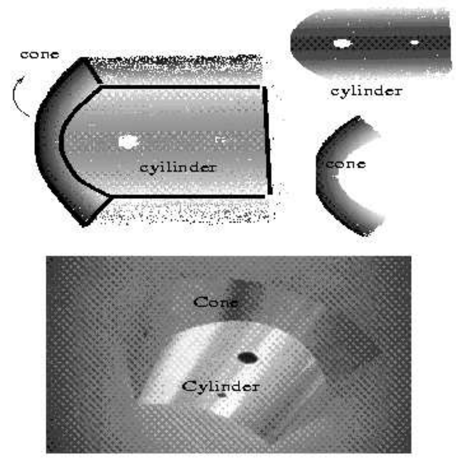

Figure 1. Constrained quadric surface recovery.

One of the assumptions underlying the work presented here is that the reverse engineering/reconstruction process is not fully automated. Computers are good at data analysis and fitting; humans are good at recognizing and classifying patterns. Thus we are working in a cooperative problem solving paradigm, where a human might hypothesize that a given relationship holds (*e.g.* two surfaces are potentially parallel), and the computer can either help verify the relationship (*e.g.* calculate the probability that they are parallel) or compute some parameter that results from the relationship (*e.g.* compute the separation between the surfaces).

From these general ideas, we have been exploring techniques to improve reverse engineering of objects from 3D point data sets. These main themes are explored in the sections that follow:

1. There are many constraints on feature relationships in manufactured objects and buildings. Exploiting these constraints improves the recovery of object models.

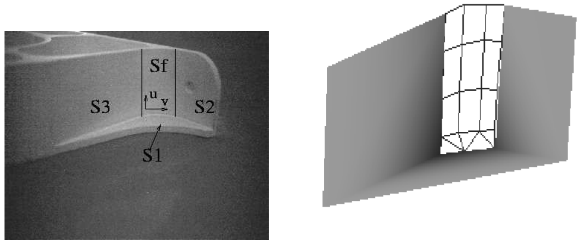

Figure 2. Constrained freeform surface recovery. Left) Object. Right) constrained planar and mesh surfaces.

2. General shape knowledge can allow recovery even when data is very noisy, sparse or incomplete.

3. Complete data acquisition can be practically impossible, but inference of much occluded data is possible.

4. Euclidean fitting is now fast enough to be practical and gives better results.

5. Many of these recovery problems require discovery of shape and position parameters that satisfy the knowledge-derived constraints. Evolutionary search methods can be used to do this search effectively.

2. Constrained reverse engineering

Parts have standard feature relationships

One of the cornerstones of the recent research in our laboratory has been constrained recovery of 3D shapes from 3D point cloud data. In this case, the constraints encode standard feature relationships such as alignment of surfaces, colinearity of features, etc. This constrained reverse engineering technique has been applied to both industrial parts and architectural scenes.

The key issue is how to incorporate design constraints into shape fitting of 3D data. Our current approach is to formulate shape fitting as constrained least-squares problem. If:

- $\vec{p}$ specifies the parameter vector for feature shapes and positions
- $\mathcal{H}$ is the least squares shape error matrix
- $C_i(\vec{p})$ are constraints over the parameters
- λ_i are penalty costs

and then minimize:

$$\vec{p}^T \mathcal{H} \vec{p} + \sum_i \lambda_i C_i(\vec{p})$$

The linear least squares error term can also be a non-linear Euclidean distance (or other) error term. This is generally a non-convex problem, so we initialize $\vec{p}$ to be the standard least-square solution and $\lambda_i = 0$. We then incrementally enforce the constraints by increasing penalty costs λ_i until the constraints are satisfied to the desired tolerances. The gradual increase ensures that the solution stays near the least-square solution and also helps avoid local minima. Experiments show that solutions initialized from different randomly perturbed starting points converge to a small cluster of nearby solutions.

We have applied this approach to engineering parts modeled by planar and quadric surfaces [17, 18]. The part shown in Figure 1 has constraints between planar, cylindrical and conical surfaces. Seven shape relation constraints were applied. All constraints can be satisfied while still maintaining close surface fitting. Applying the constraints also improves shape parameter recovery. For example, the top cylindrical surface has the true radius of 60 mm. Initial least-square quadric fitting estimated an elliptic cylinder radius of 33-46 mm. Adding the relationship constraints resulted in a circular cylinder radius estimate of 59.54 mm.

One can also apply [19] the approach to enforcing boundary constraints between freeform and quadric surfaces, while also trying to minimize surface fitting error. One application is ensuring that the freeform surface is tangential or orthogonal to a planar surface at their common boundary. Figure 2 shows three mutually orthogonal planar surfaces plus a B-spline surface tangential to two of the planes and orthogonal to the third.

More recently, we have also applied the constrained shape fitting method to architectural scenes [1]. The concepts are similar to the industrial part case as many standard architectural relationships are present, such as near perpendicularity of walls and floors, coplanarity of floors inside and outside rooms, etc. Additionally, as we know that we are recovering a building with large planar surfaces, we can recover a better model by enforcing surface flatness to displace triangle vertices onto the nearest plane. Figure 3 shows some ripples near the lower windows in the original triangulation that have been flattened.

Reconstructing models from multiple 3D point datasets requires registration of the point sets. Most registration algorithms are variants of the Iterated Closest Point (ICP) algorithm, which searches for the best corresponding points between the datasets from which the registering pose can be estimated. Our recent work [13] on pose space search has shown that one can obtain equally good registration results while avoiding local incorrect minima, from which the ICP algorithm suffers. Additionally, ICP requires a good initial estimate in order to have correct convergences, whereas our pose-space search methods allow convergence from any starting point.

Figure 3. Constrained recovery of an architectural scene. Top) Original VRML object with surface ripples most easily seen at lower left. Bottom) Flattened and constrained surfaces with fewer artifacts.

3. Knowledge-based feature extraction in noisy data

Particularly difficult problems for data-driven recovery processes are outliers, low resolution and noisy data on reflective surfaces. When we have knowledge of either the specific parts or of general design relationships that hold in a particular domain, then we can exploit this knowledge in the shape recovery process.

Boundary relationships are standardized

Figure 4 shows a surface fitting problem [12] where a cylindrical surface has a tangent join with another cylindrical surface. Data-driven surface fitting algorithms have trouble identifying a clean boundary, because surface shape variations are not distinguishable within the data noise. Using knowledge of the type of junction allows an accurate estimate of the cylinder axes, radii and intersections.

Architectural model recovery can also exploit domain knowledge. Many recent part model and building representation systems are based on triangulation models, often recovered from raw range data. These models work well with smooth surfaces, but tend to round off surface crease edges or introduce artifacts on them. We have extended the "marching triangle" surface triangulation algorithm [8] to seed triangulation at previously-located fold edges. This preserves the shape discontinuity at the edge while also allowing the accurate "decimated" triangulation of the marching triangle algorithm. Figure 5 shows part of an architectural scene with and without fold edge preservation.

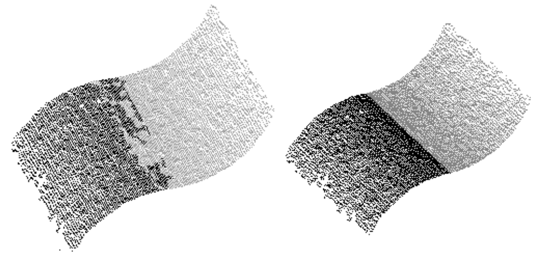

Figure 4. Cylinder/cylinder tangential surface interface with typical data-driven ragged fitting and a clean knowledge-based fitting.

Alignment relationships are standardized

Figure 6 (left) shows noisy data for a row of holes [11]. The part being reconstructed is metallic so there is a lot of surface noise from inter-reflections. In this case, as well as having a simple parametric model of the hole, we exploit additional easily obtainable knowledge about the part, namely that the holes are collinear, equally spaced and each row has equal radii holes. Using an optimization algorithm, we find the shape and position parameters that best describe the features, even with considerable noisy data.

Objects have standardized structures

We have recently applied this approach to architectural feature recovery, in this case using noisy and fragmentary 3D data [6]. Using similar optimization methods, we extract the parametric model that best fits the data fragments, as well as effectively segmenting the data by assigning appropriate 3D points to the fitted model surfaces. Figure 7 shows an example doorway fit.

4. Inference of unobservables

Constructing complete models usually requires multiple scans of an object or scene. Because of the desire to reduce acquisition costs by minimizing the number of scans while still maintaining complete coverage, researchers have

Figure 5. Building fragment with and without fold edge preservation.

Figure 6. Noisy data for a row of holes (left) and the fitted hole boundaries (right). The fitted holes are in light grey with the accepted edge points in darker grey.

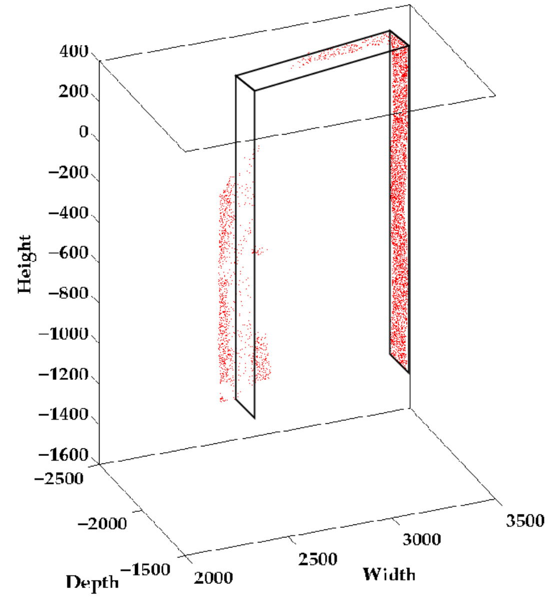

Figure 7. Noisy partial data for a doorway and the fitted parametric model.

developed view planning algorithms. From our experience with laser-based range sensors, we realized that view planning had to include a surface quality measure [7], quantifying how close the observation angle was to the surface normal at each surface point.

When we applied the view planning approach to even simple scenes [14] (See Figure 8), we found that approximately 110 views with a typical 60 degree aperture sensor were needed to observe every part of the scene. About another 100-200 were needed to observe every surface point with high accuracy. This number of scans is clearly not feasible.

The main cause of the need for so many scans is occlu-

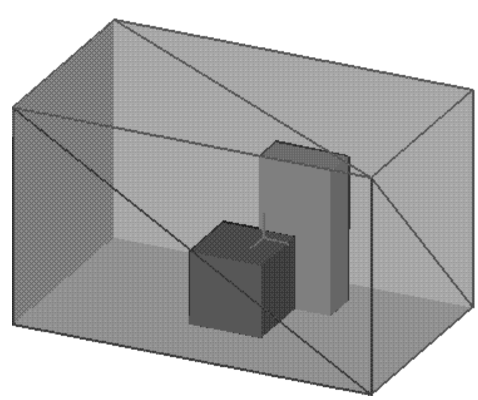

Figure 8. Simple test scene with occlusions.

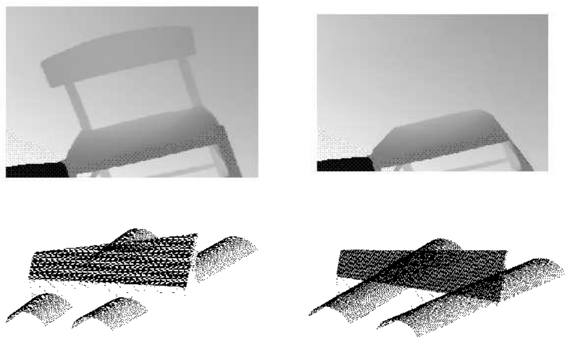

Figure 9. Original range image of scene with occluding chair back (top left) and reconstructed wall (top right). (Bottom left) two cylinders occluded by a closer surface. (Bottom right) the reconstructed cylinders.

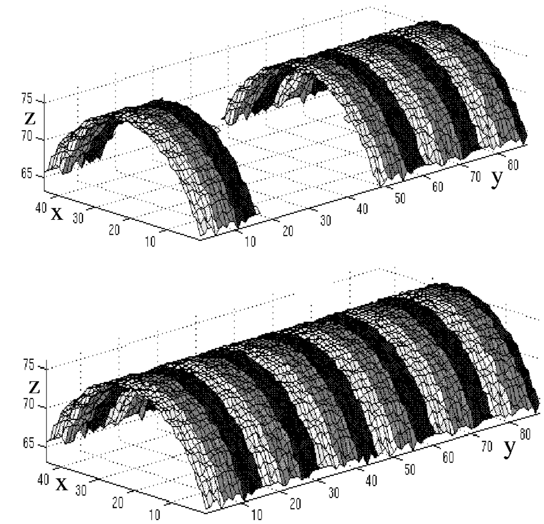

Figure 10. Cylinder before and after shape and texture restoration from behind an occlusion.

sion, where closer parts of the object or scene hide more distant parts. To obtain the missing parts, we need to position the scanner at many additional places to acquire increasingly smaller unscanned portions of the data.

Since this problem arises with even very simple parts and scenes, there probably is no "scanning" based solution to the problem. Hence, we have been investigating model and knowledge-based shape hypothesizing methods.

Standard shapes allows recovery of unobservable shape and texture

We have been recently investigating knowledge-based hypothetical reconstruction of unobserved surfaces [3, 16]. The key to reconstruction is the knowledge that the shape of the unobserved surface is usually the same as the observed portion of a surface. This allows us to project surfaces into occluded areas. As many simple surfaces have infinite extent, this requires also an estimate of the unobserved boundary [2]. We have applied this recovery process to planar and cylindrical surfaces, examples of which appear in Figure 9. Given the recovery of the surface shape, we have also been investigating [15] recovery of the surface appearance. In this case we exploit consistency of the appearance - namely either constant reflectance or repeating texture. Figure 10 shows reconstructed texture on a reconstructed cylindrical surface.

5. Better feature fitting

Euclidean distance is better and fast

One important issue in surface shape fitting and reconstruction is the choice of error metrics. For many years, the algebraic metric has been the choice for fitting quadric surfaces. If $\{\vec{x}_i\}$ is a set of 3D data points, then the algebraic fit is the A, $\vec{b}$ and c that minimizes

$$\Sigma_i(\vec{x}_i' A \vec{x}_i + \vec{b}' \vec{x}_i + c)^2$$

By the appropriate reorganization of the terms of this function, the minimization can be expressed as an eigenvalue problem with a straightforward, efficient and numerically stable solution. In the case of linear structures like planes and lines, this approach also gives the solution that minimizes the Euclidean distance to the data. Unfortunately, there is significant shape bias when fitting curved surfaces. Taubin's distance is an improvement and this can also be implemented efficiently, but still with bias. However, the Euclidean distance is usually the best:

$$\Sigma_i \mid\mid \vec{x}_i - \vec{s}_i(\vec{p}) \mid\mid^2$$

where $\vec{s}_i(\vec{p})$ is the point on the fitted surface (which is parameterized by shape and position parameters $\vec{p}$) closest to data point $\vec{x}_i$. Figure 11 shows a comparison of fittings to a real cylindrical dataset. For this important industrial shape, both the algebraic and Taubin fitting give serious errors in the fitting, while the Euclidean fit is good.

Researchers and engineers have traditionally avoided using the Euclidean distance because there is no closed form solution for general quadric surfaces thus leading to a large computational cost. (Closed forms exist for planes, elliptical cylinders and cones, which are a very practical subset of the quadric surfaces.) Recently we have reinvestigated this question because of the dramatic recent increase

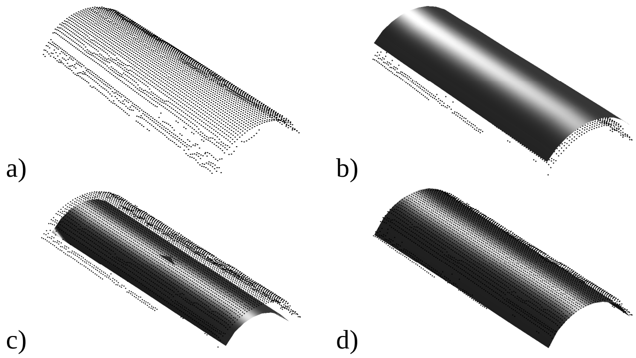

Figure 11. Comparison of three fitting algorithms to a partial cylinder. a) Original data, b) algebraic fit, c) Taubin fit, d) Euclidean fit.

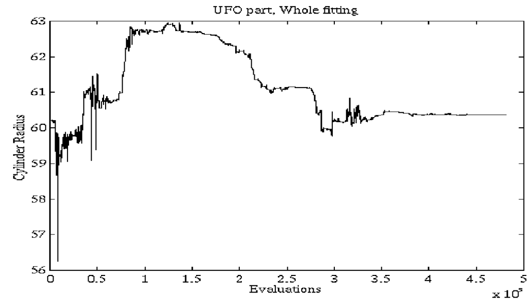

Figure 12. Convergence of radius estimate in an evolutionary constrained fitting.

in computational power [5, 4]. As well as exposing the great difference in fit quality, we have investigated the computational costs. Our efficient iterative implementation of the Euclidean fit is about 20 times slower than the closed form Taubin fit, but, in fact, the actual running time is approaching negligibility. Our implementation runs at about 3000 points/second on a 500 Mhz PC.

This implies that better quality surface fitting is now possible at reasonable costs.

6. Evolutionary structure recovery

Parameter space search to find solutions

As well as using classical optimization techniques, we have been exploring using evolutionary methods for surface fitting and 3D shape recovery [9, 10]. The advantages of evolutionary methods are: 1) Euclidean and robust error metrics are easily incorporated into the evaluation criteria and 2) initializing the optimization is not a big problem with the use of multiple "chromosomes" as the initial starting points. The main disadvantage is the larger computational cost; however, since reverse engineering a shape is usually a one-time process, the extra cost (*e.g.* a few hours rather than a few minutes) is not a problem.

The key concept to the evolutionary approach is search of the shape and position space: rather than initially finding surface and volumetric features directly from the data and then manipulating their positions, our evolutionary approach starts with the individual surface shapes (initialized by coarser segmentation processes) and manipulates their shape descriptions and positions to minimize the fitting error of all data points. In other words, the algorithm searches the space of numerical part descriptions, rather than the space of model-to-data correspondences.

Figure 12 shows how the estimates of the radius of the top cylindrical surface of the object shown in Figure 1 stabilizes as a function of number of parameter evaluations. Typical resulting inter-surface orientation error is about 0.05^o degrees.

7. Discussion and the Future

One of the issues that has arisen in the course of this research is the fragility of the reconstruction process. If reconstruction requires several stages, then: 1) the process can fail at an early stage or 2) the process can succeed, but its outputs will have results that are affected by the data errors. These 'perturbed' results then become effectively locked and affect the subsequent processes. We are exploring how to overcome the second effect and how to also reduce the failures from the first stage by looking at a one-step reconstruction process that does dataset registration, assigns point data to features, extracts feature shape parameters and accounts for standard surface shapes and constraints. Obviously this is an ambitious exploration. Optimistically, we think that the evolutionary search methods discussed above coupled with careful choices of representations will enable us to achieve and explore this goal.

We are also continuing the exploration of the knowledge-directed recovery of missing data. Many individual cases can still be investigated, but the interesting ones that we are currently exploring are 1) hypothesizing the back sides of objects based on ideas of symmetry and local space relationships and 2) recovery of unscanned 3D shape from alignment with color photographs of the unscanned areas.

Acknowledgements

This research was funded by the UK EPSRC projects GR/L25110 and GR/M97138 and EC TMR networks SMART2 (ERB FMR XCT 96 0052) and CAMERA (ERB FMR XCT 97 0127).

References

[1] H. Cantzler, forthcoming PhD thesis, Division of Informatics, University of Edinburgh, 2002.

[2] U. Castellani, S. Livatino, R. B. Fisher. “Improving Environment Modelling by Edge Occlusion Surface Completion”, Proc. Int. Symp. on 3D Data Processing Visualization and Transmission (3DPVT), Padova, Italy, forthcoming, June 2002.

[3] F. Dell’Acqua, R. Fisher, “Reconstruction of planar surfaces behind occlusions in range images”, IEEE Trans. Pat. Anal. and Mach. Intel., to appear.

[4] P. Faber, R. B. Fisher, “Euclidean Fitting Revisited”, Proc. 4th Int. Workshop on Visual Form, pp. 165-175, 28-30 May 2001, Capri, Italy. Springer-Verlag LNCS 2059.

[5] P. Faber, R. B. Fisher, “A Buyer’s Guide to Euclidean Elliptical Cylindrical and Conical Surface Fitting”, Proc. British Machine Vision Conference BMVC01, Manchester, pp 521-530, September 2001.

[6] P. Faber, R. B. Fisher, “How can we exploit typical architectural structures to improve model recovery?”, Proc. Int. Symp. on 3D Data Processing Visualization and Transmission (3DPVT), Padova, Italy, forthcoming, June 2002.

[7] N. A. Massios, R. B. Fisher. “A Best Next View Selection Algorithm Incorporating a Quality Criterion”, Proc. British Machine Vision Conference BMVC98, Southampton, pp 780–789, September 1998.

[8] N. H. McCormick, R. B. Fisher. “Edge-Constrained Marching Triangles”, Proc. Int. Symp. on 3D Data Processing Visualization and Transmission (3DPVT), Padova, Italy, forthcoming, June 2002.

[9] C.Robertson, R.B. Fisher, N. Werghi, A.P. Ashbrook. “An Evolutionary Approach to Fitting Constrained Degenerate Second Order Surfaces”, in Evolutionary Image Analysis, Signal Processing and Telecommunications, Proc. First European workshop on evolutionary computation in image analysis and signal processing (EvoIASP99). Goteborg, Sweden, pp 1-16, Springer LNCS 1596, May 1999.

[10] C. Robertson, R.B. Fisher, D. Corne, N. Werghi, A.P. Ashbrook. “Investigating Evolutionary Optimisation of Constrained Functions to Capture Shape Descriptions from Range Data”, Proc. 3rd On-line World Conference on Soft Computing (WSC3). Also in: Advances in Soft Computing - Engineering Design and Manufacturing, R. Roy, T. Furuhashi and P. K. Chawdhry (Eds), Springer, 1999, pp 455–466.

[11] C. Robertson, R. B. Fisher, N. Werghi, A. P. Ashbrook. “Fitting of Constrained Feature Models to Poor 3D Data”, Proc. Adaptive Computing in Design and Manufacture (ACDM 2000), Plymouth, UK, pp 149-160, April, 2000.

[12] C. Robertson, R. B. Fisher, “Better Surface Intersections by Constrained Evolution”, Proc. Adaptive Computing in Design and Manufacture (ACDM 2002), University of Exeter, Devon, UK, forthcoming, April 2002.

[13] C. Robertson, R. B. Fisher, “Parallel Evolutionary Registration of Range Data”, *Computer Vision and Image Understanding*, forthcoming.

[14] J. M. Sanchiz, R. B. Fisher, “A next-best-view algorithm for 3D scene recovery with 5 degrees of freedom”, Proc. British Machine Vision Conference BMVC99, Nottingham, pp 163–172, September 1999.

[15] F. Stulp. Completion of Occluded Surfaces, MSc Dissertation, University of Groningen, 2001.

[16] F. Stulp, F. Dell’Acqua, R. B. Fisher, “Reconstruction of surfaces behind occlusions in range images”, Proc. 3rd Int. Conf. on 3-D Digital Imaging and Modeling (3DIM01), Montreal, Canada, pp 232-239, June 2001.

[17] N. Werghi, R. B. Fisher, C. Robertson and A. P. Ashbrook, “Modelling Objects Having Quadric Surfaces Incorporating Geometric Constraints”, Proc. 5th Eur. Conf. on Computer Vision, Vol. II, pp 185–201, Freiburg, Germany, June 1998.

[18] N. Werghi, R. B. Fisher, C. Robertson, A. Ashbrook. “Faithful recovering of quadric surfaces from 3D range data by global fitting”, International Journal of Shape Modelling, Vol.6, No.1, pp.65-78, 2000.

[19] N.Werghi, R.B Fisher, C. Robertson, “Constrained Object Reconstruction Incorporating Free-form Surfaces”, Proc. IX Spanish Symposium on Pattern Recognition and Image Analysis, Benicssim, Spain, pp 273-280, May 2001.