



School of Informatics, University of Edinburgh

Institute of Perception, Action and Behaviour

Applying knowledge to reverse engineering problems

by

Robert Fisher

Informatics Research Report EDI-INF-RR-0209

School of Informatics
<http://www.informatics.ed.ac.uk/>

June 2004

Applying knowledge to reverse engineering problems

Robert Fisher

Informatics Research Report EDI-INF-RR-0209

SCHOOL *of* INFORMATICS

Institute of Perception, Action and Behaviour

June 2004

Final version published at Computer Aided Design, 36(6), pp 501-510, May 2004

Abstract :

This paper summarizes a series of recent research results made at Edinburgh University based on projects that apply 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 very noisy data, how to get better shape parameter estimates and how to infer data about unseen features.

Keywords : range data, reverse engineering, constraints

Copyright © 2004 by The University of Edinburgh. All Rights Reserved

The authors and the University of Edinburgh retain the right to reproduce and publish this paper for non-commercial purposes.

Permission is granted for this report to be reproduced by others for non-commercial purposes as long as this copyright notice is reprinted in full in any reproduction. Applications to make other use of the material should be addressed in the first instance to Copyright Permissions, School of Informatics, The University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, Scotland.

DRAFT: Applying knowledge to reverse engineering problems

Robert B. Fisher
School of Informatics, University of Edinburgh
rbf@inf.ed.ac.uk

Abstract

This paper summarizes a series of recent research results made at Edinburgh University based on projects that apply 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 very noisy data, how to get better shape parameter estimates and how to infer data about unseen features.

Keywords: range data, reverse engineering, constraints

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 at incorrect relative positions or artifacts arising from noisy or missing data.

For several years our research group at Edinburgh University 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, utility or engineering design. This is a “knowledge-based” approach.

We argue that exploiting this extra knowledge allows improved reverse engineering. This paper presents several different examples of the general approach, summarizing results from the full publications, which are cited within and can be found at: <http://www.dai.ed.ac.uk/homes/rbf/publications.html>.

One of the assumptions underlying the work summarized here is that the reverse engineering/reconstruction process need not be 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.* 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.
2. General shape knowledge can allow recovery even when data is very noisy, sparse or incomplete.
3. Complete data acquisition can be impossible in practice, 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. There has been much previous research on curved surface shape estimation, based either on the Euclidean distance [6] or variants of the algebraic distance [22]. Given the shape bias arising from the algebraic distance, researchers have also developed a general quadric surface extension to the algebraic distance using a gradient based weighting [52] or a shape specific approximation [31]. These fitting approaches were for single surfaces. In our case, we have used a constrained algebraic distance approach that applies shape specific constraints on all of the individual surfaces. Within the same framework, we also simultaneously apply constraints that 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 shape and 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 first term is a least squares fitting term that ensures that model surfaces lie close to the image data. The second term encodes the penalties for constraint violations. The linear least squares error term can also be a non-linear Euclidean distance (or other) error term. Minimizing this error is generally a non-convex problem, so we initialize \vec{p} to be the standard least-square solution and $\lambda_i = 0$ and then apply numerical optimization methods. We then incrementally enforce the constraints by increasing penalty costs λ_i and re-minimizing 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.

While we have only experimented with constraint functions $C()$ that use the square of the error in the constraint, one could also use a gated function that produces zero error if the constrained relationship is within a given tolerance. If this form were used, our gradient based optimization method would need to be modified as there is a discontinuity at the tolerance point. One possible approach is to use evolutionary methods mentioned in Section 6. Then the constraint can be simply ignored in the evaluation function if the specified tolerance is satisfied.

We have applied this approach to engineering parts modeled by planar and quadric surfaces [55, 56]. The surfaces are extracted from range images or point clouds by a segmentation process based on 1) shape discontinuity detection, 2) boundary constrained noise smoothing, 3) principal curvature based local shape classification and finally 4) quadric surface fitting. A recent comparison in [26] concluded that in many ways our single image planar surface segmentation algorithm had the best performance among current algorithms.

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 the approach [57] to enforcing inter-surface boundary constraints between freeform and quadric surfaces, while also trying to minimize surface fitting error. This differs from the pure constrained freeform

surface case [30] and the pure quadric surface case [55, 56]. We also considered constraints between non-adjacent surfaces as well as connectivity constraints. We satisfied the constraints effectively exactly using a numerical optimization process instead of an equation-solving approach [44], using the data projection method of [33]. 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 applied the constrained shape fitting method to architectural scenes [9]. 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 by displacing triangle vertices onto the nearest plane. Figure 3 shows some ripples near the lower windows in the original triangulation that have been flattened. The data for this example was acquired by an expensive range sensor, but some of our other examples [9] have used sparse 3D triangulated point sets obtained from structure-and-motion recovery from video sequences. Due to the sparseness of the 3D triangulated data features, we needed a different segmentation process to assign vertices to planar surface patches. After that, constrained surface adjustment and fitting proceeded in the same way as the part shape recovery. The use of structure-and-motion data would probably not be so useful in the other techniques presented in the paper as that data tends to be quite sparse and much noisier than range sensor data.

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 [42] 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 greatly (*e.g.* [5, 12] and many improvements). Additionally, ICP requires good initial estimates in order to have correct convergence, whereas our pose-space search methods allow convergence from any starting point.

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 [41] 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 interface and its parameters, such as the cylinder axes, radii and intersections in this example.

Architectural model recovery can also exploit domain knowledge. Many recent part model and building representation systems are based on triangulation models [45], 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 [32] the “marching triangle” surface triangulation and multiple surface fusion algorithm [24, 25] to seed triangulation [11] at previously-located fold edges (using RANSAC [20]). This preserves the shape discontinuity at the fold edges while also allowing the accurate “decimated” triangulation of the marching triangle algorithm. Figure 5 shows part of an architectural scene without and with fold edge preservation.

Alignment relationships are standardized

Figure 6 (left) shows noisy data for a row of holes [40]. The part being reconstructed is metallic so there is a lot of surface noise in the shape data arising from inter-reflections. In this case, as well as having a simple parametric model of the hole, we exploit additional user supplied 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 point cloud data [18]. While one could use constrained feature space search methods (*e.g.* [1, 21]), here we use a constrained optimization method in which the constraints are built into the optimization process (*e.g.* [6]), to fit parametric shape models (*e.g.* [8, 13, 46]). While constraints are not well exploited in reverse-engineering [54], and often features are extracted independently (*e.g.* [4]), here we simultaneously fit, establish point-to-feature correspondences and estimate parameters. Using similar optimization methods as above, we extract the position and shape of a 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, where the doorway has 6 positional and 3 shape degrees of freedom.

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 developed view planning algorithms. From our experience with laser-based range sensors, we realized that view planning had to include a surface quality measure [35], 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 [43] (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 (unless a wide-field of view panoramic sensor is used [29]). The main cause of the need for so many scans is occlusion, 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. While there has been much previous work on viewplanning (*e.g.* [37, 51]), this work dealt with simple nearly convex objects, and so did not encounter the problems arising from having many inter-part or object occlusions.

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 [14, 50]. There is work on recognizing objects from range data, considering occlusions [36], but here we are attempting to recover from them. 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 [23, 34]. 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 [10]. 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 [49] recovery of the surface appearance [15]. 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 (*e.g.* [4]). If $\{\vec{x}_i\}$ is a set of 3D data points, then the algebraic fit [22] is the A , \vec{b} and c that minimizes

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

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 and other variants [52, 53, 7, 58, 28] or shape specific methods [31] are improvements and can be implemented efficiently, but still with bias. However, the Euclidean distance is usually the best:

$$\sum_i \|\vec{x}_i - \vec{s}_i(\vec{p})\|^2$$

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 recent dramatic increase in computational power [17, 16]. 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 [38, 39]. 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 that arises in parameter space search instead of parametric surface growing in data space, either in 2D [4] or 3D [19], or for triangulated 3D surfaces [27]. 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. Simpler parts like that in Figure 1, with about 20 constraints, required about 30 minutes computation on a 200 Mhz Sun workstation, which is probably equivalent to about 5 minutes on a current PC. As the number of parameters grows, the computation time will grow, in part from the additional terms in the evaluation function, but also minima will be harder to find. On the other hand, the minimization always starts with good feature position estimates coming from the feature fitting, so the the parameter vector is always close to the optimal. Hence, progress can be more rapid. Further, this approach is an “anytime” algorithm, meaning it can be stopped at any time with a feasible, if suboptimal, solution.

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° degrees.

7. Discussion and the Future

This paper is largely a summary of recent research results at Edinburgh University. There are two other significant streams of research that should be compared to that presented here. The first is research from the University of Utah [48, 47], who have also been investigating constrained reconstruction of parts from 3D data sets, particularly parts with pocket profiles. They categorized the types of engineering knowledge as domain specific and pragmatic (roughly corresponding to our use of a restricted set of shapes), and functional constraints

(corresponding to our use of interfeature constraints). They exploit this knowledge to select surface types and manufacturing actions. Thus, with some user assistance, planar features that bound pockets are found. The contour that is swept to form the pocket can then be found automatically. Shape and positional constraints are represented and solved in a manner similar to our approach. A particularly notable aspect of their research that we have not addressed is the automatic inference of feature relationships and how these can be used to reduce the degrees of freedom of the reconstruction by equating parameters (which we have exploited [38]).

The other significant research stream is that at the University of Cardiff [2, 3]. That research has also followed a route similar to that in this paper, exploiting designed-in relationships to improve reconstruction. In their case, a much larger set of relationships were explored, and the constraints arising from the relationships were used to reduce the dimensionality of the reconstruction parameter space. A sequential numerical constraintment processes was used, which allowed them to detect (which we also did) and automatically reject (which we did not) inconsistent constraints. A nice alternative to fitting tangential and blend surfaces was to parameterize swept two dimensional features, with the cross-section of the inter-surface join/blend as the two-dimensional feature.

What all three of these projects have in common is an appreciation of the role of intent in the design of artifacts, and how this intent is expressed in relationships that can be exploited in the reverse engineering process. There are differences in the numerical optimization process, but all express both the geometric fitting and the shape constraints in a numerical evaluation function that can be effectively solved to reconstruct the constrained shape.

While this paper is more of a summary paper, in addition to the commonality discussed above, the paper also points to several other pieces of research not in common with the others, namely: beautification by constrained triangulation flattening, application to architecture as well as parts, constrained fitting of spline surfaces, the benefits of Euclidean fitting, the practical impossibility of complete scene scanning, shape and texture hypothesising, triangulation with fold edge preservation, and higher level reverse engineering by using structures parameterized at the object level rather than the feature level.

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 explore and achieve 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). Many thanks to all of the researchers who contributed to this sequence of research.

References

- [1] R. Anderl, R. Mendegen, “Modelling with constraint: Theoretical foundation and application”, CAD, 28(3), pp 155-168, 1996.
- [2] P. Benkő, R. R. Martin, T. Várady Algorithms for Reverse Engineering Boundary Representation Models Computer Aided Design 33 (11) 839-851, 2001.
- [3] P. Benkő, G. K/’os, T. Várady, L. Andor, R. R. Martin Constrained Fitting in Reverse Engineering Computer Aided Geometric Design, 19, 173-205, 2002.

- [4] P. J. Besl, “Analysis and Interpretation of Range Images”, Springer, Berlin-Heidelberg-New York, 1990.
- [5] P. J. Besl, N. D. McKay, “A method for registration of 3D shapes” *IEEE Trans. Pat. Anal and Mach. Intel.*, 14(2), pp 239-256, 1992.
- [6] R. M. Bolle, D. B. Cooper. “On optimally combining pieces of information, with application to estimating 3-D complex-object position from range data”, *IEEE Trans. Pat. Anal and Mach. Intel.*, 8(5), pp 619-638, Sept 1986.
- [7] F. L. Bookstein, “Fitting conic sections to scattered data”, *Comp. Graph. and Image Proc.*, 9, pp 56-71, 1979.
- [8] C. Brechbuehler, G. Gerig, O. Keubler, “Parameterisation of closed surfaces for 3-d shape description”, *Comp. Vis. and Image Under.*, 61(2), pp 154-170, 1995.
- [9] H. Cantzler, R. B. Fisher, M. Devy, “Improving architectural 3D reconstruction by plane and edge constraining”, *Proc. British Machine Vision Conf.*, Cardiff, pp 43-52, 2002.
- [10] 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, pp 672-675, June 2002.
- [11] X. Chen, F. Schmidt. “Surface Modelling of range data by constrained triangulation”, *Computer-Aided Design*, 26(8), pp 632-645, 1999.
- [12] Y. Chen, G. Medioni, “Object modeling by registration of multiple range images”, *Object Modeling by Registration of Multiple Range Images, Image and Vision Comp.*, 10(3), pp. 145-155, 1992.
- [13] G. Danuser, M. Stricker, Parametric model fitting: from inlier characterization to outlier detection. *IEEE Trans. Pat. Anal and Mach. Intel.*, 20(3), pp 263-280, 1998.
- [14] F. Dell’Acqua, R. Fisher, “Reconstruction of planar surfaces behind occlusions in range images”, *IEEE Trans. Pat. Anal. and Mach. Intel.*, 24(4), pp 569-575, April 2002.
- [15] A. A. Efros, K. Leung. “Texture synthesis by non-parametric sampling”, *Proc. Int. Conf. on Comp. Vis.*, pp 1033-1038, 1999.
- [16] 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.
- [17] 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.
- [18] 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, pp 824-833, June 2002.
- [19] R. B. Fisher, A. 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.
- [20] M. A. Fishler, R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Tech report 213, AI Center, SRI International, 1980.
- [21] E. Grimson, “The role of geometric constraints”, MIT Press, London, 1990.
- [22] K. T. Gunnarsson, F. B. Prinz. “CAD model-based localization of parts in manufacturing”, *IEEE Computer*, 20(8), pp 66-74, Aug 1987.
- [23] L. Herault, R. Horaud. “Figure-ground discrimination: a combinatorial optimization approach”, *IEEE Trans. Pat. Anal and Mach. Intel.*, 9(15), pp 899-914, 1993.

- [24] A. Hilton, A. Stoddart, J. Illingworth, T. Winderatt. “Marching Triangles: Range Image Fusion for Complex Object Modelling”, Proc. Int. Conf. on Image Proc, vol. 2, pp. 381–384, 1996.
- [25] A. Hilton, A. Stoddart, J. Illingworth, T. Winderatt. “Implicit surface-based geometric fusion”, Comp. Vis. and Image Under., 69(3), pp 273-291, March 1998.
- [26] 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.
- [27] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, W. Stuetzle, “Surface reconstruction from unorganized points”, Comp. Graphics, 26(2), pp 71-78, 1992.
- [28] K. Kanatani, “Renormalization for biased estimation”, Proc. Int. Conf. Comp. Vis., pp 599-606, 1993.
- [29] Klein K., Sequeira V. “The View-Cube: An Efficient Method of View Planning for 3D Modelling from Range Data”, 5th IEEE Workshop on Applications of Computer Vision (WACV’2000), Palm Springs (CA), USA, December 2000.
- [30] J. P. Kruth, A. Kersrens. “Reverse engineering modelling of free-form surfaces from point clouds subject to boundary conditions”, J. of Materials Processing Technology, 76, pp 120-127, 1998.
- [31] G. Lukacs, A. D. Marshall, R. R. Martin. “Faithful least-square fitting of spheres, cylinders, cones and tori for reliable segmentation”, Proc Eur. Conf. on Computer Vision, Vol 1, pp 671-686, June 1998.
- [32] N. H. McCormick, R. B. Fisher. “Edge-Constrained Marching Triangles”, Proc. Int. Symp. on 3D Data Processing Visualization and Transmission (3DPVT), Padova, Italy, pp 348-351, June 2002.
- [33] W. Ma, J. P. Kruth, “Parameterization of randomly measured points for least squares fitting of B-spline curves and surfaces”, Computer-Aided Design, 27(9), pp 663-675, 1995.
- [34] S. Mahamud, K. K. Thornber, L. R. Williams. “Segmentation of salient closed contours from real images”, Proc. Int. Conf. on Comp. Vision., pp 891-897, 1999.
- [35] 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.
- [36] P. G. Mulgahonkar, C. K. Cowan, J. DeCurtins. “Understanding object configurations using range images”, IEEE Trans. Pat. Anal. and Mach. Intel., 14(2), pp 303-307, 1992.
- [37] M. K. Reed, P. K. Allen, I. Stamos. “Automated model acquisition from range images with view planning”, Proc. Int. Conf. on Comp. Vis. and Pat. Rec., pp 72-77, 1997.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] C. Robertson, R. B. Fisher, “Better Surface Intersections by Constrained Evolution”, Proc. Adaptive Computing in Design and Manufacture (ACDM 2002), Ed. I.C. Parmee, University of Exeter, Devon, UK, Springer, pp 133-142, April 2002.

- [42] C. Robertson, R. B. Fisher, "Parallel Evolutionary Registration of Range Data", *Computer Vision and Image Understanding*, 87, pp 39-50, 2002.
- [43] 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.
- [44] B. Sarkar, C. H. Menq. "Parameter optimization in approximating curves and surfaces to measurement data", *Computer Aided Geometric Design*, 8, pp 267-290, 1991.
- [45] Sequeira V., Gonçalves J.G.M. "3D Reality Modelling: Photo-Realistic 3D Models of Real World Scenes", Proc 1st International Symposium on 3D Data Processing Visualization and Transmission (3DPVT 2002), Padova, Italy, pp 776-783, June 19-21, 2002.
- [46] L. Staib, J. S. Duncan, "Boundary finding with parametrically deformable models", *IEEE Trans. Pat. Anal and Mach. Intel.*, 14(11), 1061-1075, 1991.
- [47] W. B. Thompson, J. C. Owen, H. J. de St. Germain, S. R. Stark, T. C. Henderson, "Feature-Based Reverse Engineering of Mechanical Parts", *IEEE Trans. Robotics and Automation*, 15(1): 57-66, 1999.
- [48] H. J. de St. Germain, S. R. Stark, W. B. Thompson, T. C. Henderson, "Constraint Optimization and Feature-Based Model Construction for Reverse Engineering", Proc. ARPA Image Understanding Workshop, 1997.
- [49] F. Stulp, "Completion of Occluded Surfaces", MSc Dissertation, University of Groningen, 2001.
- [50] 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.
- [51] K. A. Tarabanis, R. Y. Tsai, A. Kaul. "Computing Occlusion-free viewpoints", *IEEE Trans. Pat. Anal and Mach. Intel.*, 18(3), pp 279-292, 1996.
- [52] G. Taubin. "Estimation of planar curves, surfaces and non-planar space curves defined by implicit equations with applications to edge and range image segmentation", *IEEE Trans. Pat. Anal and Mach. Intel.*, 13(11), pp 1115-1138, Nov 1991.
- [53] G. Taubin, "An improved algorithm for algebraic curve and surface fitting", Proc. Int. Conf. Comp. Vis., pp 658-665, 1993.
- [54] T. Varady, R. R. Martin, J. Cox, "Reverse engineering of geometric models - an introduction", *CAD*, 29(4), pp 255-269, 1997.
- [55] 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.
- [56] 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.
- [57] 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.
- [58] Z. Zhang, "Parameter estimation techniques: a tutorial with application to conic fitting", *Image and Vision Comp.*, 15, pp 59-76, 1997.




Figure 1. Constrained quadric surface recovery.





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



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.




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



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




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

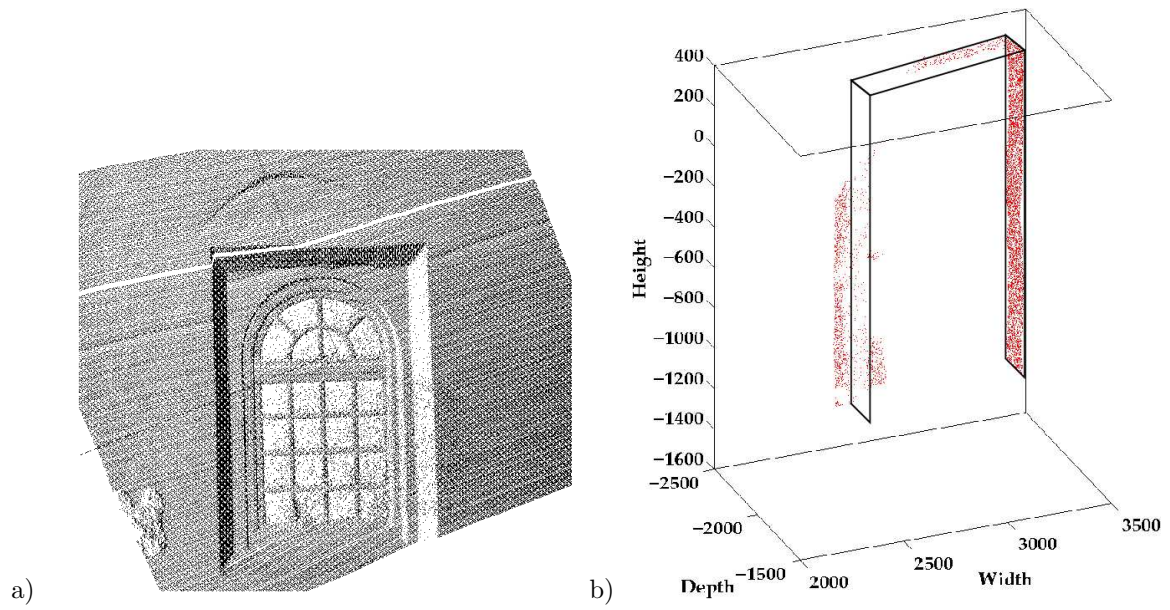


Figure 7. a) Noisy incomplete data for a doorway, b) the fitted parametric model and selected data.




Figure 8. Simple test scene with two interior occluding objects.




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.

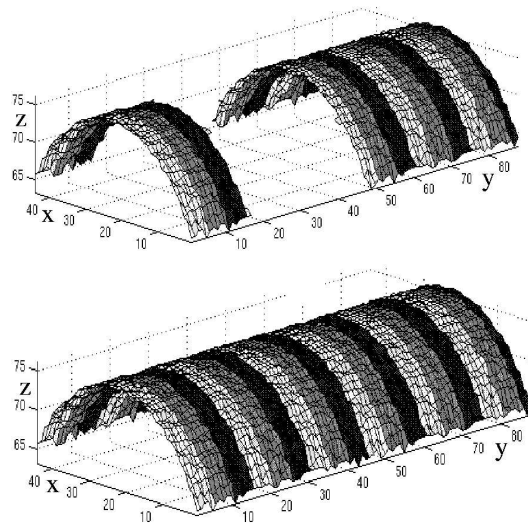


Figure 10. Cylinder before and after shape and texture restoration from behind an occlusion.




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.