Solving architectural modelling problems using knowledge

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 applying domain knowledge of standard shapes and relationships to solve or improve architectural reconstruction problems. The problems considered are how to enforce known relationships when data £tting, how to extract features even in very noisy data, how to get better shape parameter estimates and how to infer data about unseen features.

1. Introduction

Traditional processes for reconstructing buildings from 3D datasets have been initially data (*e.g.* triangulated models) and parametric surface (*e.g.* quadric surface) driven. These approaches have been successful, 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 and speci£c architectural knowledge. The process is not "model-based" reconstruction 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 buildings that humans usually construct are not arbitrary: their shapes 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 architectural reconstruction. 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 architectural reconstruction process need not be fully automated. Computers are good at data analysis and £tting; humans are good at recognizing and classifying patterns. Thus, we work in a cooperative problem solving paradigm, where a human might hypothesize that a given relationship holds (*e.g.* two walls 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 walls).

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

1. There are many constraints on feature relationships in buildings. Exploiting these constraints improves the recovery of 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 or otherwise missing data is possible.

4. 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 building reconstruction

Buildings have standard feature relations

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 [5] or variants of the algebraic distance [21]. 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 [47] or a shape speci£c approximation [27]. These £tting approaches were for single surfaces. In our case, we have used a constrained algebraic distance approach that applies shape speci£c 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.

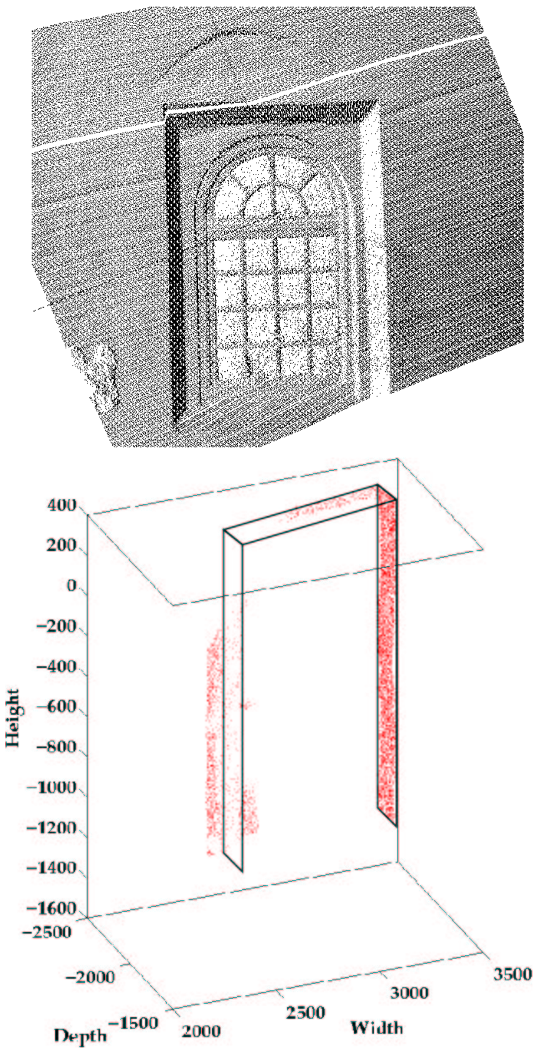

Figure 1. (top) Noisy incomplete data for a doorway, (bottom) the £tted parametric model and selected data.

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

- $\vec{p}$ speci£es 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 £rst term is a least squares £tting 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 satis£ed 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. A gated form would be particularly appropriate for architecture, as buildings always deviate somewhat from their design, either through construction variations, subsequent modi£cation or subsidence. If a gated form were used, our gradient based optimization method would need to be modi£ed as there is a discontinuity at the tolerance point. One possible approach is to use the evolutionary methods mentioned in Section 6. Then the constraint can be simply ignored in the evaluation function if the speci£ed tolerance is satis£ed.

We have applied the constrained shape £tting method to architectural scenes [7, 8], where many standard architectural relationships are present, such as near perpendicularity of walls and ¤oors, coplanarity of ¤oors inside and outside rooms, etc. We also considered constraints between non-adjacent surfaces as well as connectivity constraints. We satis£ed the constraints effectively using a numerical optimization process instead of an equation-solving approach [40], using the data projection method of [29]. While these buildings only have constrained planar surfaces, we have

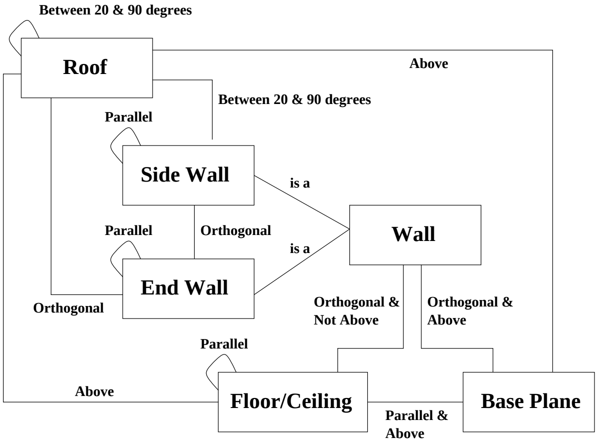

Figure 2. Semantic net encoding typical architectural relationships.

also developed techniques for constraints involving quadric and freeform surfaces [49, 50].

The constraints that are applied in a reconstruction can arise from two processes. The easiest is by a reconstruction engineer specifying the constraints. More interestingly, constraints could be hypothesized automatically. At a purely data-driven level, one could apply statistical tests to validate hypothesized relationships (*e.g.* perpendicularity of surfaces, as in [2]). We have investigated using higher level architectural knowledge to supplement the local feature relationships [8, 7]. The knowledge is encoded in a semantic net, see Figure 2. The reconstruction process attempts to label scene surface and edge features as instances of the given building parts whenever the data feature relationships satisfy the model relationships. Assignment of labels uses a search algorithm tolerant to shape and position errors. Once data features are labeled, the ideal constraints relating the features in the semantic net are used to constrain the feature positions during reconstruction.

3. Knowledge-based shape improvement

Feature shapes allow local improvements

As we know that we are recovering buildings with large planar surfaces, we can recover better models by enforcing surface ¤atness through displacing triangle vertices onto the nearest plane [8, 7]. Figure 9 shows some ripples near the lower windows in the original triangulation that have been ¤attened. The data for this example was acquired by an expensive range sensor, but some of our other examples [7] have used sparser 3D triangulated point sets obtained from structure-and-motion recovery from video sequences. Figure 3 shows some surfaces before and after ¤attening. Notice that ¤attening does not occur everywhere, but instead tries to improve only the coplanar features, so that the windows, door and bicycle are preserved. 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 £tting 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 this paper as that data tends to be quite sparse and much noisier than range sensor data.

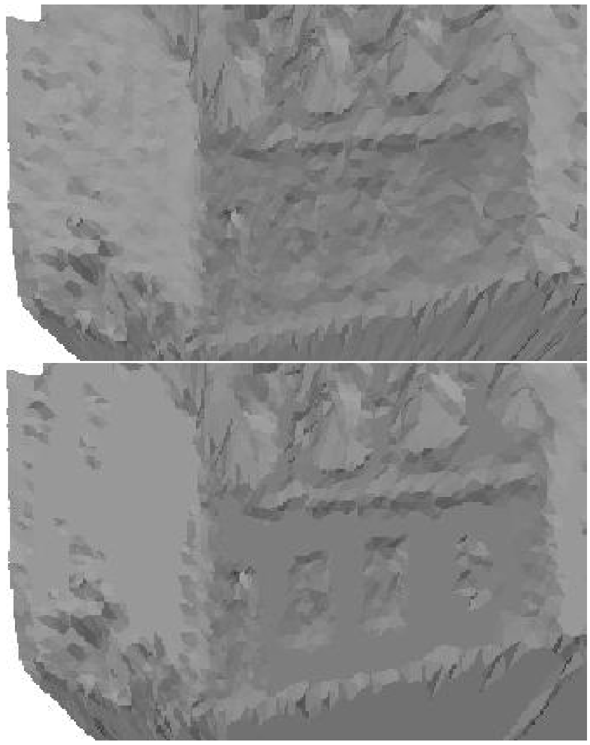

Figure 3. (top) Raw triangulated model. (Bottom) Flattened model.

Many recent part model and building representation systems are based on triangulation models [41], 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 [28] the "marching triangle" surface triangulation and multiple surface fusion algorithm [23, 24] to seed triangulation [10] at previously-located fold edges (using RANSAC [19]). This preserves the shape discontinuity at the fold edges while also allowing the accurate "decimated" triangulation of the marching triangle algorithm. Figure 4 shows part of an architectural scene without and with fold edge preservation.

4. Knowledge-based feature extraction in noisy data

Figure 4. Building fragment without and with fold edge preservation.

Particularly dif£cult problems for data-driven recovery processes are outliers, low resolution and noisy data on re¤ective surfaces. When we have knowledge of either the speci£c buildings or of general design relationships that hold in a particular domain, then we can exploit this knowledge in the shape recovery process.

Buildings have simple edge relations

Figure 5 (top) shows sets of potential fold edges extracted from a triangulated mesh. Knowing that most buildings have edges aligned perpendicularly allows us to extract the three principal directions in the architecture. This information allows the edges and the linking surfaces to be recti£ed to full orthogonality. Figure 5 (bottom) shows the hypothesized principal directions extracted from the full edge set.

Buildings have standardized structures

We have recently applied a parametric model approach to architectural feature recovery, in this case using noisy and fragmentary 3D point cloud data [17]. While one could use constrained feature space search methods (*e.g.* [1, 20]), here we use a constrained optimization method in which the constraints are built into the optimization process (*e.g.* [5]), to £t parametric shape models (*e.g.* [6, 12, 42]). While constraints are not well exploited [48], and often features are extracted independently (*e.g.* [3]), here we simultaneously £t, 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 £ts the data fragments, as well as effectively segmenting the data by assigning appropriate 3D points to the £tted model surfaces. Figure 1 shows an example doorway £t, where the doorway has 6 positional and 3 shape degrees of freedom.

5. Inference of unobservables

Constructing complete models usually requires multiple scans of a 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 [31], 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 [39] (see Figure 6), 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 £eld of view panoramic sensor is used [26]). The main cause of the need for so many scans is occlusion, where closer parts of the 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 view planning (*e.g.* [35, 46]), that work dealt with simple nearly convex objects, and so did not encounter the problems arising from having many occlusions. Since this occlusion problem arises with even very simple scenes, there probably is no “scanning” based solution to the problem.

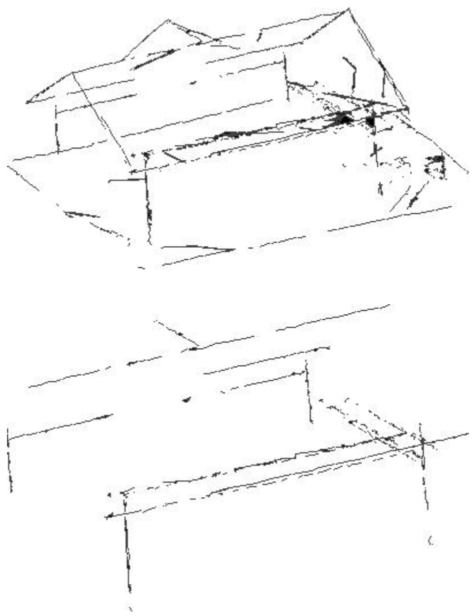

Figure 5. Top) extracted fold edges, bottom) hypothesized orthogonal principal directions.

Standard shapes allow recovery of unobservable shape and texture

We have been investigating knowledge-based hypothetical reconstruction of unobserved surfaces [14, 45]. There is work on recognizing objects from range data, considering occlusions [33], 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 [22, 30]. This allows us to project surfaces into occluded areas. As many simple surfaces have in£nite extent, this requires also an estimate of the unobserved boundary [9]. We have applied this recovery process to planar and cylindrical surfaces, an example of which appears in Figure 7.

As well as occluded forward facing structures, we have also investigated hypothesizing the back-facing sides of columns [32]. Figure 8 (top) shows a subset of a triangulated scan of the prayer hall of the Edinburgh Central Mosque. The model is formed from the fusion of 11 cylindrical scans, each with about 12 million 3D points, but still many portions of the model are incomplete. Here, we see some columns that are missing part or all of 1, or 3 sides. The close up in Figure 8 (mid) shows some partial columns with most of 3 sides, plus some holes and Figure 8 (bottom) shows the extended columns.

Given the recovery of the surface shape, we have also been investigating [44] recovery of the surface appearance [16]. In this case we exploit consistency of the appearance - namely either constant re¤ectance or repeating texture.

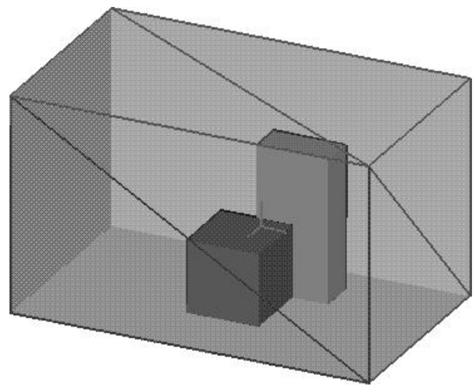

Figure 6. Simple test scene with two interior occluding objects.

6. Evolutionary structure recovery

Parameter space search to £nd solutions

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 [38] 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.* [4, 11] 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.

Besides using classical optimization techniques, we have been exploring using evolutionary methods for surface £tting and 3D shape recovery [36, 37]. The key concept to the evolutionary approach is search of the shape and position space: rather than initially £nding 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 £tting 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.

The advantages of evolutionary methods are: 1) Euclidean and robust error metrics are easily incorporated into the evaluation criteria and 2) initializing the optimization

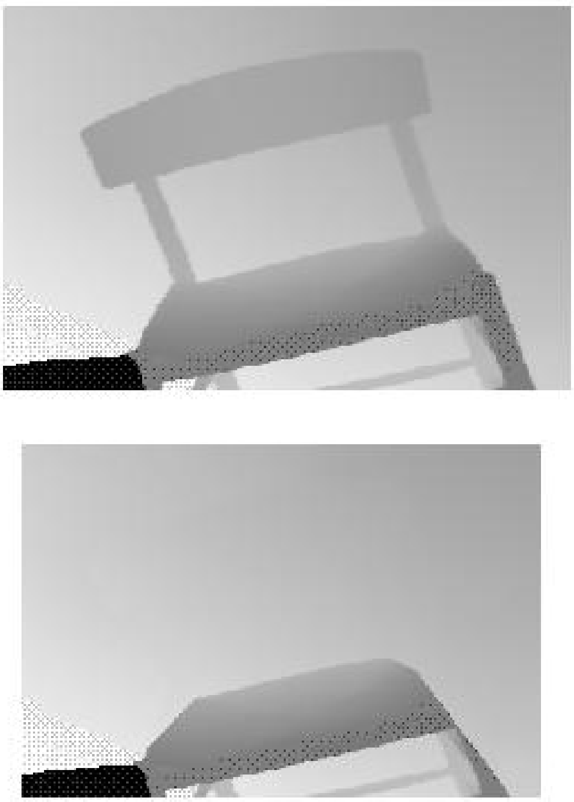

Figure 7. Original range image of scene with occluding chair back (top) and reconstructed wall (bottom).

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 [3] or 3D [18], or for triangulated 3D surfaces [25]. However, since reconstruction is usually a one-time process, the extra cost (*e.g.* a few hours rather than a few minutes) is not a problem. Simpler reconstructions 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 £nd. On the other hand, the minimization always starts with good feature position estimates coming from the least square feature £tting, so 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.

7. Discussion and the Future

This paper is largely a summary of recent research results at Edinburgh University. Many of the techniques described here heve been applied to industrial part reverse engineering as well as building reconstruction. Our research in that direction as been in¤uenced by the excellent research at the Univ. of Utah [43] and Univ. of Cardiff [2]. We have also been much in¤uenced by the intensity image-based architectural reconstruction at Berkeley [13] and Cambridge University [15], the video sequence analysis of Pollefeys [34] and the range image analysis at the EC Joint Research Centre at Ispra [41].

What these projects have in common is an appreciation of the role of intent in the design of structures, and how this intent is expressed in relationships that can be exploited in the reconstruction process.

While this paper is more of a summary paper, in addition to the commonalities with the approaches mentioned above, the paper also points to several other pieces of research not in common with the others, namely: the practical impossibility of complete scene scanning, occluded shape hypothesising, beauti£cation by constrained triangulation ¤attening, triangulation with fold edge preservation and higher level reconstruction 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 £rst 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.

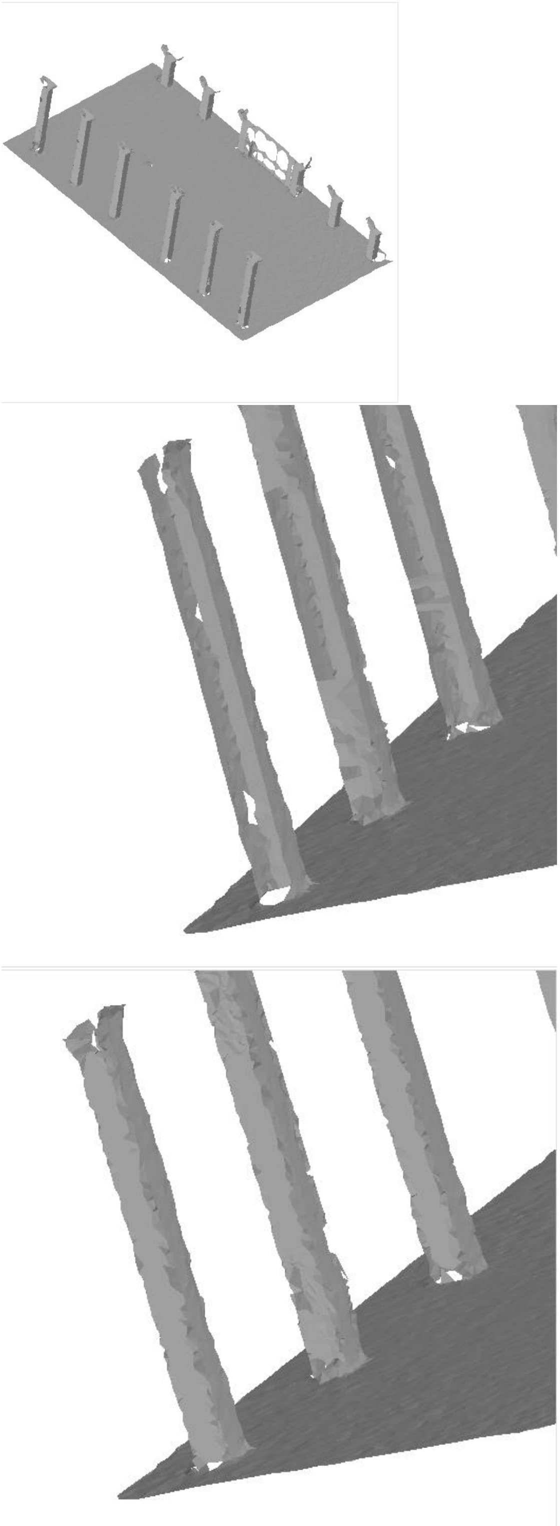

Figure 8. (top) Partial model of mosque. (middle) Close-up showing incomplete columns. (bottom) Close-up showing completed columns.

Acknowledgements

This research was funded by the UK EPSRC projects GR/L25110, GR/M00166 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. Benkö, R. R. Martin, T. Várady Algorithms for Reverse Engineering Boundary Representation Models Computer Aided Design 33 (11) 839-851, 2001.

[3] P. J. Besl, “Analysis and Interpretation of Range Images”, Springer, Berlin-Heidelberg-New York, 1990.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] H. Cantzler. “Improving architectural 3D reconstruction by constrained modelling”, PhD Thesis, School of Informatics, Univ. of Edinburgh, 2003.

[9] 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.

[10] X. Chen, F. Schmidt. “Surface modelling of range data by constrained triangulation”, Computer-Aided Design, 26(8), pp 632-645, 1999.

[11] 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.

[12] G. Danuser, M. Stricker, Parametric model £tting: from inlier characterization to outlier detection. IEEE Trans. Pat. Anal and Mach. Intel., 20(3), pp 263-280, 1998.

[13] P. E. Debevec. “Modeling and Rendering Architecture from Photographs”, Ph. D.Thesis, UC Berkeley, 1996.

[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. R. Dick, P. H. S. Torr, R. Cipolla. “A Bayesian Estimation of Building Shape Using MCMC”, Proc. Eur. Conf. on Comp. Vis, Part II, pp 852, 2002.

[16] A. A. Efros, K. Leung. "Texture synthesis by non-parametric sampling", Proc. Int. Conf. on Comp. Vis., pp 1033-1038, 1999.

[17] 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.

[18] 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.

[19] M. A. Fishler, R. C. Bolles. Random sample consensus: a paradigm for model £tting with applications to image analysis and automated cartography. Tech report 213, AI Center, SRI International, 1980.

[20] E. Grimson, "The role of geometric constraints", MIT Press, London, 1990.

[21] K. T. Gunnarsson, F. B. Prinz. "CAD model-based localization of parts in manufacturing", IEEE Computer, 20(8), pp 66-74, Aug 1987.

[22] L. Herault, R. Horaud. "Figure-ground discrimination: a combinatorial optimization approach", IEEE Trans. Pat. Anal and Mach. Intel., 9(15), pp 899-914, 1993.

[23] A. Hilton, A. Stoddart, J. Illingworth, T. Windeatt. " Marching Triangles: Range Image Fusion for Complex Object Modelling", Proc. Int. Conf. on Image Proc, vol. 2, pp. 381–384, 1996.

[24] A. Hilton, A. Stoddart, J. Illingworth, T. Windeatt. "Implicit surface-based geometric fusion", Comp. Vis. and Image Under., 69(3), pp 273-291, March 1998.

[25] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, W. Stuetzle, "Surface reconstruction from unorganized points", Comp. Graphics, 26(2), pp 71-78, 1992.

[26] K. Klein, V. Sequeira. "The View-Cube: An Ef£cient 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.

[27] G. Lukacs, A. D. Marshall, R. R. Martin. "Faithful least-square £tting of spheres, cylinders, cones and tori for reliable segmentation", Proc Eur. Conf. on Computer Vision, Vol 1, pp 671-686, June 1998.

[28] 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.

[29] W. Ma, J. P. Kruth, "Parameterization of randomly measured points for least squares £tting of B-spline curves and surfaces", Computer-Aided Design, 27(9), pp 663-675, 1995.

[30] 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.

[31] 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.

[32] G. Miller, "Reconstruction of the unobserved sides of building columns", undergraduate dissertation, School of Informatics, Univ. of Edinburgh, 2003.

[33] P. G. Mulgahonkar, C. K. Cowan, J. DeCurtins. "Understanding object con£gurations using range images", IEEE Trans. Pat. Anal. and Mach. Intel., 14(2), pp 303-307, 1992.

[34] M. Pollefeys, R. Koch, L. Van Gool. "Self-Calibration and Metric Reconstruction Inspite of Varying and Unknown Intrinsic Camera Parameters", Int. J. Comp. Vis, Vol 32, No. 1, August 1999, pp. 7-25.

[35] 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.

[36] 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.

[37] 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", In: Advances in Soft Computing - Engineering Design and Manufacturing, R. Roy, T. Furuhashi and P. K. Chawdhry (Eds), Springer, 1999, pp 455–466.

[38] C. Robertson, R. B. Fisher, "Parallel Evolutionary Registration of Range Data", *Computer Vision and Image Understanding*, 87, pp 39-50, 2002.

[39] 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.

[40] B. Sarkar, C. H. Menq. "Parameter optimization in approximating curves and surfaces to measurement data", Computer Aided Geometric Design, 8, pp 267-290, 1991.

[41] V. Sequeira, J. G. M. Gonçalves. "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.

[42] L. Staib, J. S. Duncan, "Boundary £nding with parametrically deformable models", IEEE Trans. Pat. Anal and Mach. Intel., 14(11), 1061-1075, 1991.

[43] 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.

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

[45] 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.

[46] 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.

[47] G. Taubin. "Estimation of planar curves, surfaces and non-planar space curves de£ned 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.

[48] T. Varady, R. R. Martin, J. Cox, “Reverse engineering of geometric models - an introduction”, CAD, 29(4), pp 255-269, 1997.

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

[50] 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.

Figure 9. 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.