Self-Calibrating Surface Reconstruction for the ModelMaker

S. J. Cunnington and A. J. Stoddart
Centre for Vision, Speech and Signal Processing
University of Surrey, Guildford, Surrey GU2 5XH, UK
`[a.stoddart|s.cunnington]@ee.surrey.ac.uk`

Abstract

The ModelMaker [8] is a commercially available hand-held laser scanner mounted on an articulated arm. In this paper we present results from post-processing of ModelMaker data and a novel technique for recalibrating the data while postprocessing. We demonstrate significant noise reduction of the resulting surface data. These results suggest that lower specification, possibly cheaper arms, may be used to get results as good as those currently obtained with high specification arms.

1 Introduction

Traditional range scanners are mounted on XY platforms or scan objects mounted on rotating platforms. These scanners have the limitation that they cannot capture complex geometry in a single scan.

Recently 3D Scanners [8] has produced a sensor called ModelMaker that facilitates the capture of more complex geometries. The sensor is based on a small hand-held laser striper mounted on an articulated arm that measures in real time the position and orientation of the striper. R. B. Fisher et al [2] presented a prototype based on a similar concept in 1996 . Using this approach a cloud of point measurements can be captured. The data is subsequently processed into a single surface using a surface fusion algorithm developed at Surrey [4, 5] and incorporated into the ModelMaker product.

The fusion algorithm is based on a volumetric scheme in which an implicit surface representation is created and subsequently triangulated using a marching cubes algorithm. In this paper we present results for a postprocessing scheme based on surface refitting, surface decimation and data recalibration.

The surface refitting scheme takes the output surface and the original point data and performs an optimisation with respect to a conventional cost function. This is based on the idea that the volumetric scheme, because it uses voxels of finite size can lose accuracy relative to the point data (although this is implementation dependent). On the other hand it is recognised that the volumetric scheme is very good at extracting the correct topology from complex surfaces, which is a weakness of surface based schemes.

The ModelMaker output surface is a product of the marching cubes algorithm and can have a large number of faces. It is desirable to reduce the number of faces and we therefore use a decimation scheme based on edge collapse. Both of these algorithms are derived

BMVC 1998 doi:10.5244/C.12.79

from the work of Hoppe [6], but we present quantitative results for the ModelMaker sensor.

Finally we introduce a novel technique for recalibrating the striper pose during surface reconstruction. Taken together all these techniques significantly reduce the noise present in the output surface.

Figure 1: The ModelMaker sensor

In sections 2 and 3 we develop theory relevant to the sensor and the surface fitting approach. In section 4 we motivate and describe our approach, followed by results in section 5 and a conclusion.

2 Theory: ModelMaker

In this section we develop a theoretical model of the ModelMaker measurement process. We suppose that there are $j = 1..N$ "stripes" of data produced, each with $i = 1..M_j$ data points denoted $\vec{x}^S_{j,i}$. These are 3D coordinates in the striper reference frame. We assume that the covariance of these data points is given by the 3x3 matrix W_S. We will in fact make a crude isotropic noise assumption, i.e. that the covariance matrix is diagonal and that the variance is σ^2_S for each of the 3 components.

At the instant stripe j is measured the arm reports a position $\vec{t_j}$ and a rotation matrix R_j. We follow the notation of Pennec [7] and concatenate these into a pose f_j. The rotation matrix can also be expressed as a unit vector n and a rotation θ in radians, or a single vector $p = \theta n$. We will use the notation

- Application of $f = (R, t)$ to $x : y = f * x = Rx + t$.
- Composition of $f_1 = (R_1, t_1)$ with $f_2 = (R_2, t_2)$:
  $f = f_2 \circ f_1 = (R_2 R_1, R_2 t_1 + t_2)$

Thus in the world coordinate system the measured point is given by

$$\vec{x}^W_{j,i} = f_j * \vec{x}^S_{j,i} \tag{1}$$

Following the approach of Pennec to modelling frame noise we should write the noisy frame measurement $\hat{f}_j$ in terms of the true (unknown) frame f_j and a right error frame e, i.e. $\hat{f}_j = f_j \circ e$. For a motivation of this formulation of frame error see Pennec. Briefly this formulation allows the formulation of a 6x6 frame covariance matrix from e if the errors are small and the rotation is expressed as $p = \theta n$.

The true arm pose f_j can be decomposed into the end-effector pose f_j^E and the striper alignment f^A. The latter must be computed by a calibration procedure, and is subject to error. In summary

$$\hat{f}_j = f_j \circ e = f_j^E \circ f^A \circ e_j \tag{2}$$

3 Theory: Surface Fitting

Our surface representation is a set of triangles with a list of vertices denoted V and topological connectivity of these vertices denoted by K, thus $S = \{V, K\}$ in the notation of [6]. The set of point measurements is denoted by $X = \{x_1..x_D\}$. The surface fitting process may be formulated as a minimisation problem over the surface vertex positions $V = \{v_1..v_n\}$ and mesh topology with an objective function given by

$$E_{tot}(K, V) = E_{dist}(K, V) + E_{rep}(K) + E_{reg}(K, V) \tag{3}$$

The objective function is a sum of a data fidelity term measuring the sum of the squared distances from data to the nearest point on the surface.

$$E_{dist}(K, V) = \sum_{i=1}^{D} d^2(x_i, S) \tag{4}$$

a penalty term based on the number of vertices

$$E_{rep}(K) = \lambda_{rep} N_K \tag{5}$$

and a regulariser

$$E_{reg}(K, V) = \lambda_{reg} \sum_{\{j,k\} \in K} |v_j - v_k|^2 \tag{6}$$

The regulariser constant λ_{reg} need only play a role when there are triangles with no associated data points, after sufficient decimation it may be ignored. The representation constant λ_{rep} may be chosen by the user to make a compromise between detail and size of the representation.

4 Motivation

In this section we discuss the errors that arise in the ModelMaker measurement process. Firstly we explain how the point measurements are processed into an output surface. ModelMaker is used to “paint” patches of the object surface. A foot pedal starts and stops the acquisition of what we term “micropatches”. A series of points on adjacent stripes are triangulated to form the micropatch. A threshold ensures that holes or step edges are not closed. The micropatches are then fused using a volumetric fusion technique [4] into a resulting surface.

There are potentially many sources of error in this process, but we will only consider what we believe are the 3 most significant. Firstly the measured striper points are noisy with some rms σ_S. The striper pose will be in error by e_j, which will cause the point rms in the world coordinate system to be a larger value σ_W. Finally the volumetric fusion technique may cause additional errors, depending on the exact method chosen. These errors may arise if each voxel stores a condensed version of the voxel contents or if the polygonisation algorithm makes simplifying assumptions. More sophisticated approaches [5, 3] can circumvent this, but there is a time vs accuracy trade-off.

We propose the following post-processing chain which begins with the output from the ModelMaker. We use the fused surface and the original micropatch data. The surface becomes an initial guess $S_0 = \{K_0, V_0\}$ and the micropatch data is converted to a point set $X = \{x_{j,i|j=1..N,i=1..M_j}\}$.

We firstly minimise $E(K, V)$ with respect to V using S_0 as a starting point. This step has the effect of correcting some of the errors caused by the volumetric fusion, but benefits from having a good starting point. The next step is to choose some value for λ_{rep} and optimise over both K, V so as to decimate the surface.

How do we assess the effect of these operations? One useful measure is the rms distance of data points from the closest point on the surface, which we denote d_{rms}. The effect of step 1 is always to reduce d_{rms}, provided that the regularising constant is small. The more we decimate the more we will inevitably raise d_{rms} once again.

We now come to the question of "self-calibration". The points in each stripe are converted to the world coordinate system by a rigid body transform $f_j^E \circ f^A \circ e_j$ where we model the error by e_j.

Many of the stripes overlap so there is redundant information present and we argue that this information may be used to recalibrate the data. We propose a simple way of doing this by a (small) transform g_j applied to each stripe. We can reformulate the distance cost as

$$E_{dist}(K, V, G) = \sum_{i=1}^{D} d^2(g_j * x_{j,i}, S) \tag{7}$$

To minimise E_{tot} simultaneously over V, K and G presents formidable obstacles, so we introduce some simplifications. Firstly we iterate over alternate minimisations with respect to V and $G = \{g_1..g_N\}$. This breaks the problem into two familiar tasks, namely surface refitting (optimisation over V) and registration of several point sets to a surface (optimisation over G). The "recalibration" step is nothing more than the registration typically solved using the Iterated Closest Point algorithm [1, 10].

The second simplification is to group sets of stripes together into micropatches and let j denote not stripes but micropatches. This considerably reduces the number of new degrees of freedom at the risk of not fully being able to recalibrate individual stripes. However from our results we are clearly able to retain much of the benefit. This is probably due to the fact that some significant sources of error vary over the arm workspace slowly, and the short term random noise on our arm is not too large in comparison.

We have implemented the above optimisations using standard techniques published in the literature, mainly [6, 1]. Space considerations do not permit us to reproduce these algorithms in detail.

5 Results

All of the data presented in this section was collected on a "ModelMaker L" at the University of Surrey. The striper is mounted on a Faro Bronze B06 arm. The arm has a quoted accuracy of 0.15mm single point repeatability (one standard deviation) for point measurements. The orientation accuracy is not stated. The more expensive Faro Silver arm has a spec of 0.04mm. One source of error is the temporal update rate which is 60Hz for the Bronze, whereas the Silver can be synchronized to within 0.1ms.

What is the accuracy of the points measured by ModelMaker? 3D Scanners provide spherical and cubic calibration objects and software routines to compute the errors. A portion of the calibration sphere (radius = 37.925mm) was scanned with a single micropatch and with several overlapping micropatches. A sphere was fitted to these measurements and the rms error e_{rms} from the sphere was computed. For one micropatch the rms error was 0.1278mm and for several micropatches the rms error was 0.2387mm. [There are 4 degrees of freedom and several hundred points so we believe that the discrepancy is not accounted for by overfitting.]

From these results we conclude that striper point noise in the striper coordinate system, σ_S, is about 0.1278mm or less, but that the error in the world coordinate system can rise to 0.2387mm and above due to arm errors. [3D Scanners achieve less than 0.1mm rms for the striper in known position.] These figures are consistent with the notion that there is no benefit in either sensor being very much more accurate than the other!

In the following discussion we will use two measures of surface quality. The more reliable is the sphere rms error e_{rms} which is about as good a measure of the instrument accuracy as can be easily got. The second is the rms distance to surface d_{rms} which is a questionable measure of accuracy since in the limit of enough triangles this can always be reduced to zero. However when the number of measurements per face rises above 5, i.e. about 10 points per vertex it is apparent that we are *not overfitting*. [Typically there are approximately twice as many faces as vertices.] In this case we argue that d_{rms} is a measure of "local accuracy", i.e. the accuracy of points relative to other nearby points. This accuracy measure will not be sensitive to global distortions but will reliably quantify measurement of local shape variation. In the absence of a large set of calibration objects d_{rms} is useful.

5.1 Sphere

We now present results for data taken from the calibration sphere. All surfaces are shown with flat-shaded rendering for easier interpretation.

Figure 2(a) shows the original mesh S_0 (4861 faces) constructed from the 12498 points obtained via ModelMaker. The accuracy of the original point data is $e_{rms} = 0.2277$. Six micropatches were used in the surface reconstruction.

After the initial refit stage, registration of the micropatches is then performed. These newly registered points are then used in the next refit stage. This process is then repeated. The results are summarised in Table 1.

The d_{rms} for the original surface, as shown in Table 1, was 1.0692. This is worse than e_{rms} reflecting approximations made in the surface fusion and possibly some boundary effects. The fusion voxel size was set at 2mm. After doing a surface refit, d_{rms} dropped from 1.0692 to 0.1715. The original number of triangles were 4861, and after decimation

Table 1: Results From Sphere Refinement

Stage	No. of Triangles	d_{rms} mm	e_{rms} mm
Initial 1	4861	1.0692	0.2277
Refit 1	4861	0.1715	
Decimation 1	876	0.1844	
Initial 2	876	0.1572	0.1717
Refit 2	876	0.1389	
Decimation 2	735	0.1419	
Initial 3	735	0.1390	0.1655
Refit 3	735	0.1352	
Decimation 3	720	0.1355	
Initial 4	720	0.1351	0.1640
Refit 4	720	0.1331	
Decimation 4	715	0.1332	

they were reduced to 876. d_{rms} rose slightly to 0.1844 as expected.

Once registration of the micropatches had been performed, the newly registered point set was used in the 2nd surface refit, where d_{rms} then dropped from 0.1572 (improvement due to registration) to 0.1389. The rest of the results can be seen from Table 1.

At 17 data points per face we are not overfitting. The initial refitting causes a big drop in d_{rms} drop from 1.0692 to 0.1715. The recalibrating stage then causes a further improvement by a factor 0.68 to 0.1332. This suggests that it is of real value.

More compelling evidence for real improvement is provided by comparing the e_{rms} values on the point data. The drop is from 0.2277 to 0.1640, i.e. improvement by a factor 0.72. We conclude that we have reduced the errors that can be ascribed to the arm significantly, possibly by as much as half.

5.2 Corner

Figure 2(b) shows the original mesh constructed from the 27957 points obtained via ModelMaker. Eight micropatches were used in the surface reconstruction, one of which is shown in figure 2(c). Figure 2(d) shows the final result obtained and in Figure 2(e) shows the same result including the triangulation.

d_{rms} initially starts at 0.3864 and then drops to 0.1120 after refitting. After decimation, which has reduced the number of triangles from 9332 down to 53, d_{rms} increases to 0.1528, which can be seen in Table 2. Once registration of the micropatches has been done, d_{rms} drops to 0.1382, and after the second refit the d_{rms} drops to 0.1148. The refit improves by a factor 0.36 and the recalibration step by a further factor 0.75

5.3 Gravy Dish

Finally we show results for a real object. The object is a porcelain dish with a patterned relief. On average each part of the surface is scanned twice. In figure 3 (a) we show the fusion output with $d_{rms} = 0.3936$ and 12236 faces. There are 10 micropatches and

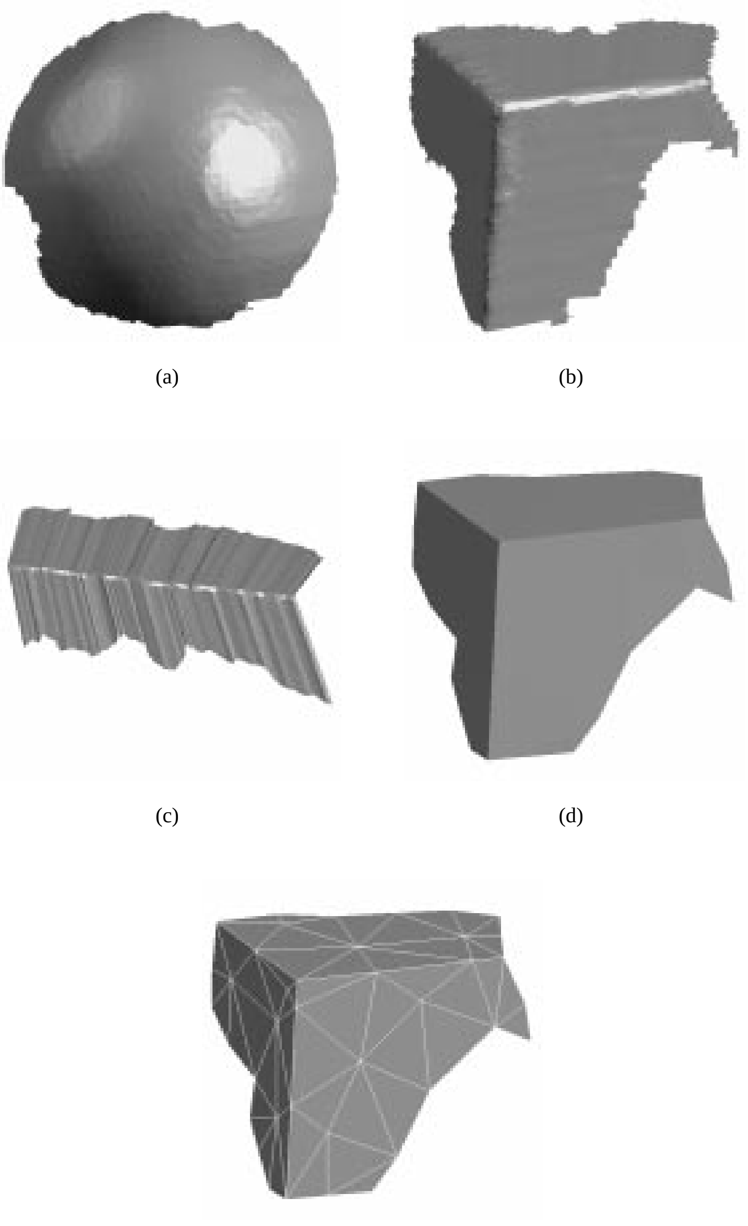

(a)

(b)

(c)

(d)

(e)

Figure 2: (a) Sphere. (b) Corner. (c) A micropatch. (d) Final result. (e) Final result showing triangulation.

Table 2: Results From Corner Refinement

Stage	No. of Triangles	d_{rms}
Initial 1	9332	0.3864
Refit 1	9332	0.1120
Decimation 1	53	0.1527
Initial 2	53	0.1382
Refit 2	53	0.1148

a total of 80095 points. In figure 3 (b) we show a slightly decimated and refitted result ($d_{rms} = 0.1503$) which can be seen in Table 3. There are 23.6615 points per face. In figure 3 (c) we show the results after recalibration ($d_{rms} = 0.0964$).

Table 3: Results From Dish Refinement

Stage	No. of Triangles	d_{rms}
Initial 1	12214	0.3936
Refit 1	12214	0.1446
Decimation 1	3817	0.1503
Initial 2	3817	0.1142
Refit 2	3817	0.0980
Decimation 2	3359	0.1006
Initial 3	3359	0.1022
Refit 3	3359	0.0958
Decimation 3	3290	0.0964

5.4 Toy Man

The final result is that of a Toy Man, (in fact a footballer well known for his lachrymose tendencies). The model head is only 40mm high and we ran the fusion software at a voxel size of 2mm. The ModelMaker has a workspace of about half a meter so the very small model is a stringent test of its operational limits. If the head were part of a much larger scene the voxel size would be realistic.

In summary we have chosen an object that we expect to visibly highlight the improvements we aim for. The initial mesh has 3699 triangles and 39471 points. It has $d_{rms} = 0.9029$. After refitting we reduce d_{rms} to 0.3440 (3699 tris) and recalibration and refitting reduces it still further to $d_{rms} = 0.1349$ (tris= 1535, 22 pts per face).

The results are shown in figure 3.

6 Conclusion

We have presented a postprocessing chain for ModelMaker data and shown qualitative and quantitative results that show significant improvement. The resulting surface reduces the

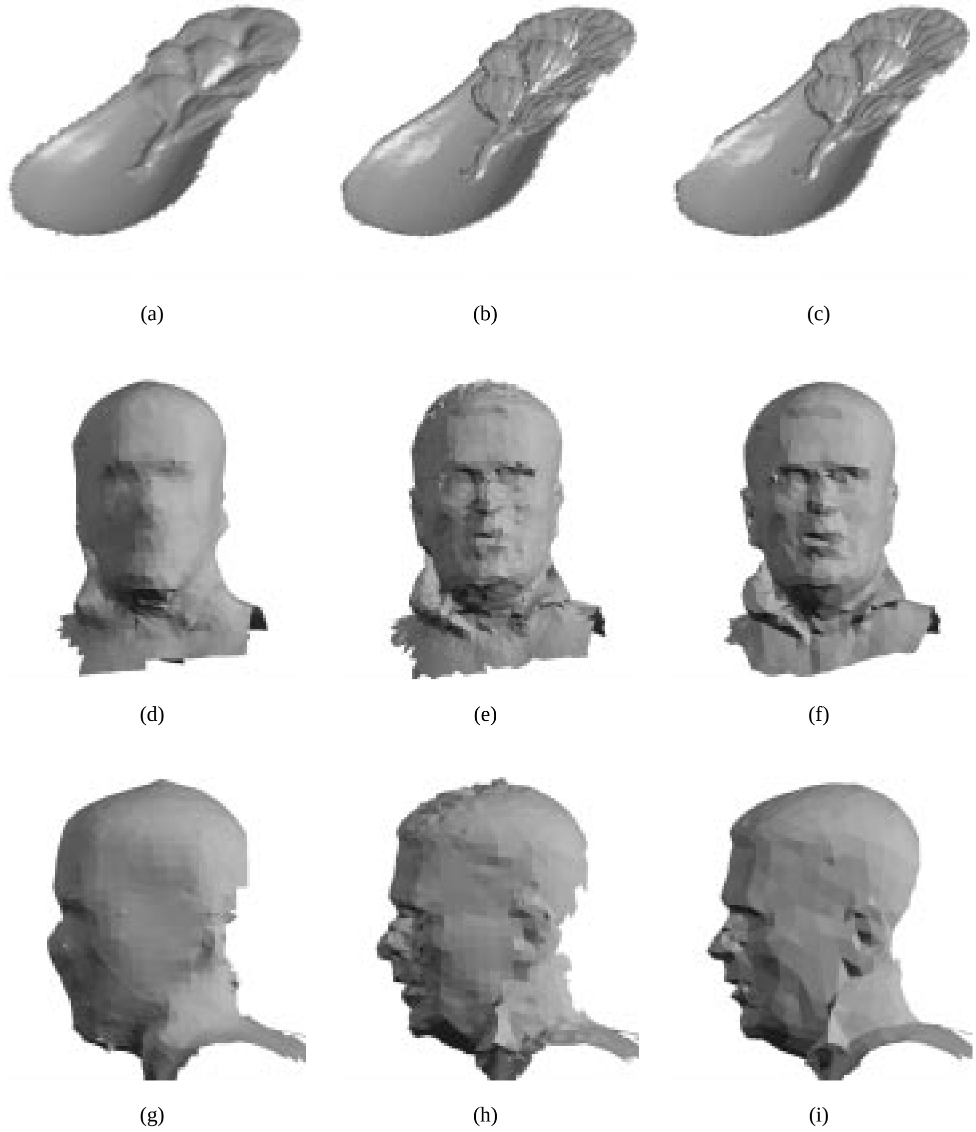

Figure 3: (a) The fused result from the micropatches. (b) Result after first refit and decimation (no registration). (c) Result from third refit. (d) Original constructed mesh (front view). (e) First refit (front view). (f) Final refit with registration (front view). (g) Original constructed mesh (side view). (h) First refit (side view). (i) Final refit with registration (side view).

errors caused in fusion and by arm errors so that we may approach the limiting accuracy of the striper.

It is a commercial reality that the higher the cost of the arm the better the accuracy that may be obtained. Arm prices vary from £1000 to £60000 and alternative non-mechanical technologies are available. An intriguing prospect of this work is the possibility of using cheaper arms to obtain results as good as those currently obtained by expensive arms.

Future work should include examining a stripe based self calibration scheme to see how it compares with the micropatch scheme. It would also be of interest to see how the recalibration scheme works with spline based surface representations [9].

We would like to acknowledge helpful comments from Peter Champ and Juri Radkov of 3D Scanners and Adrian Hilton. The work was funded by EPSRC and Copernicus project VIRTUOUS.

References

[1] P.J. Besl and N.D. McKay. A method for registration of 3D shapes. *IEEE Trans. Pattern Analysis and Machine Intell.*, 14(2):239–256, 1992.

[2] R. B. Fisher, A. Fitzgibbon, A. Gionis, M. Wright, and D. Eggert. A hand-held optical surface scanner for environmental modelling and virtual reality. In *Virtual Reality World*, Stuttgart, Germany, 1996.

[3] A. Hilton, A. J. Stoddart, J. Illingworth, and T. Windeatt. Marching triangles: range image fusion for complex object modelling. In *1996 Int. Conference on Image Processing*, pages II381–384, Lausanne, Switzerland, 1996.

[4] A. Hilton, A. J. Stoddart, J. Illingworth, and T. Windeatt. Reliable surface reconstruction from multiple range images. In *Fourth European Conference on Computer Vision*, pages 117–126, Cambridge, U.K., 1996.

[5] A. Hilton, A. J. Stoddart, J. Illingworth, and T. Windeatt. Implicit surface based geometric fusion. *Computer Vision and Image Understanding*, 69(3):273–291, 1998.

[6] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. In *SIGGRAPH*, pages 19–25, 1993.

[7] X. Pennec and J.-P. Thirion. A framework for uncertainty and validation of 3D registration methods based on points and frames. *to appear in International Journal of Computer Vision*, 1997.

[8] 3D Scanners. `http://www.3dscanners.com/`. South Bank Technopark, 90 London Road, London SE1 6LN, U.K.

[9] A. J. Stoddart and M. S. Baker. Reconstruction of smooth surfaces with arbitrary topology adaptive splines. In *Fifth European Conference on Computer Vision*, pages II 241–254, Freiburg, Germany, 1998.

[10] A. J. Stoddart, S. Lemke, A. Hilton, and T. Renn. Estimating pose uncertainty for surface registration. *Image and Vision Computing*, 16:111–120, 1998.