Extracting Axially Symmetric Geometry From Limited 3D Range Data

Andrew Willis †, Xavier Orriols ‡, Senem Velipasalar †,
Xavier Binefa ‡, David B. Cooper †
† Brown University, Providence, RI 02912
‡ Computer Vision Center (UAB), 08193 Bellaterra, Spain

December 6, 2000

Abstract

This paper describes a new, novel low computational-cost approach for recovering geometric structure and pose information of a 3D surface rotationally symmetric about an axis given an unorganized set of 3D range data which covers only a small patch of the surface. Self occluding contours are not seen here. All the parameters necessary to completely define the corresponding surface are estimated. The estimated parameters of the data patch consist of a line describing the axis and a set of parameters which describe the profile curve with respect to the axis. Global algebraic surfaces and local splines consisting of frustrums of cones are used as models for the surface estimation.

1. Introduction

With great strides in recent technology of 3D scanners [1] and 3D surface reconstruction from a moving camera [10], there is great interest in 3D object surface representation and the inference of significant 3D structure from noisy clouds of unorganized data. Whereas it is not difficult to estimate a smooth interpolation using B-splines, unconstrained algebraic surfaces, Fourier series, or other representation, the challenge is to estimate structure which is important for interpretation, recognition, or manipulation, and this usually means *constrained* representations. Among these are the quadric surfaces (spheres, cylinders, cones, planes) and the generalized cylinders (a cross sectional curve swept along an axis). Among the generalized cylinders, the simplest and most frequently encountered is the circular cross sectional generalized cylinder where the circle plane is perpendicular to a straight axis and circle radius varies as the circle is swept along the axis. In many applications, the entire generalized cylinder is seen within an image and the circle radius can be reliably estimated from the self occluding contours (the silhouette) in the image. In other applications, this will not work, either because the self occluding contours are not seen or because the camera is very close to the generalized cylinder and the occluding contours in a single image are insufficient for determining the radius. In this paper, we deal with estimation of the generalized cylinder when the occluding contours are not seen and the sensed data consists of x, y, z points measured by a 3D scanner, a Vitana ShapeGrabber [1] in our case. This is then the problem of estimating a surface which is rotationally symmetric about an axis from a patch of noisy measurement data of the surface. When the patch is small, estimating the cylinder axis and radius function for the patch is a formidable challenge. There are 3 considerations which make this seemingly simple problem a significant challenge:

1. Estimation of a small data-patch axis. If the patch is small, the axis orientation is not apparent, and often estimation results in two or even three solutions which must be resolved by additional information.

2. Estimating the shape of the profile curve (see section 3.1) is not difficult if the data-patch axis can be estimated. But estimating the distance of the profile curve from the axis is equivalent to estimation of the surface curvature in the direction perpendicular to the axis, and this is a challenge when the data patch is small.

3. For greatest estimation accuracy, the objective function minimization for 3D surface model estimation is nonlinear numerical minimization. The objective function is highly multimodal, i.e., has many local minima. Getting a computationally fast estimate of the global minimum is the challenge here.

Among the applications where this situation arises are: estimation of the 3D surface shape from a small fragment of a pot that has been made on a potter's wheel for archaeology applications such as reconstruction of a mathematical model of a pot from fragments; estimation of a manufactured or architectural 3D structure when the structure is only a portion of the generalized cylinder, i.e., the structure is that portion of a generalized cylinder that lies on one side of a plane which slices through the cylinder parallel to its

axis; or the structure is a portion of a generalized cylinder where the self occluding contours are not present in an image or 3D measurement set because of occlusion by other objects.

1.1. Related Work

Much work has been dedicated to the general problem of inferring object pose and shape from 3D range data. In [3], it is pointed out that generalized cylinders are effective models for shape recognition and representation. A classic reference on the use of generalized cones for shape representation may be found in [5]. Nevatia *et. al.* have introduced methods of estimating two categories of generalized cylinders, using only their self-occluding contours from sequences of images [6]. In [4], straight homogeneous generalized cylinders, (SHGCs), are estimated by extracting object contours and applying a shape-from-shading technique. In contrast to both of these approaches, models estimated here are very specific, i.e. straight homogeneous generalized cylinders with circular cross section, and the self-occluding contours are assumed to be missing. Potmann *et. al.* have used concepts developed in algebraic geometry, Plücker coordinates, to provide a novel linear solution to estimating axially symmetric geometry from 3D range data in [7]. However, this method does not use an objective function that makes full use of the information in the data, and the resulting estimates are not as accurate as possible. This is a serious drawback if the data is noisy or the object surface is not perfectly axially symmetric which is the case in most applications of interest. We use their solution as a starting estimate in our approach.

2. Axially Symmetric Surfaces

A surface, rotationally symmetric about an axis, which we refer to as "axially symmetric," is a surface that can be parametrized as

$$X(r(z),\theta,z) = (r(z)\cos\theta,\ r(z)\sin\theta,\ z) \qquad (1)$$

In the literature, this is a special case of a Simple Homogeneous Generalized Cylinder, SHGC, one which has circular cross section. Here the surface axis is the z-axis, $r(z)$ is the radius function – it varies as a function of z – and θ takes all values between 0 and 2π. The graph of $r(z)$ is, in the literature, referred to as the object *outline*, the *generating contour*, or the *profile curve*. In (1), the radius function $r(z)$ is a single-valued function of z. In this paper, we generalize (1) and consider profile curve (r,z) where r may take multiple values for one or more values of z. We treat two representations for $r(z)$, one being a linear spline (i.e. a sequence of straight-line segments), and the other being a planar algebraic curve (i.e. a planar implicit polynomial curve). Note, in the former case, the surface representation is a sequence of frustrums of cones, whereas in the latter case the surface is an algebraic surface constrained to be axially symmetric. The problem of interest is: given an unorganized set of 3D measured points of a small patch of a larger axially symmetric surface, estimate a surface geometry model for the patch. *We see that this geometry is completely specified by an axis in 3D and a profile curve with respect to that axis.*

3. Models and Objective Functions

The pose of a rotationally symmetric object is defined up to a sign by the axis of symmetry. We have parametrized the axis of symmetry using the standard parametric equation of a line:

$$l = \vec{v}t + \vec{q}_o$$

This equation requires 5 parameters, 3 which specify a point $\vec{q}_o$, and 2 which define the unit vector $\vec{v}$ in the direction of the line by specifying 2 angles in spherical polar coordinates. Taking the axis of symmetry to be the z-axis in an r-z plane, the profile curve is specified as a curve in this plane. We study two models for a profile curve.

1. A piecewise frustrums of cones approximation of the 3D surface. This model uses a continuous sequence of straight line segments to represent the profile curve. For L line segments, the complete profile curve is then given by:

$$r(z) = \bigcup_{i=1}^{L} c_i z + w_i \qquad (2)$$

   Here the z-axis is partitioned into L intervals, a straight line having parameters (c_i, w_i), specifies the line in the i^{th} interval, and adjacent line segments meet at the interval boundary points (see Fig. 3). If $r(z)$ takes more than one value for some values of z, a parametrization and partition as a function of arc length is used.

2. A global implicit polynomial approximation of the 3D surface. This model uses a single implicit polynomial curve to represent the profile curve. The resulting IP profile curve model of degree n takes the following form:

$$f_n(r,z) = \sum_{0\le j+k\le n;\ j,k\ge 0} a_{jk} r^j z^k = 0 \qquad (3)$$

   Expanding this form, we have:

$$a_{00} + a_{10}r + a_{01}z + \ldots + a_{n0}r^n + a_{(n-1)1}r^{n-1}z + \ldots + a_{0n}r^n = 0$$

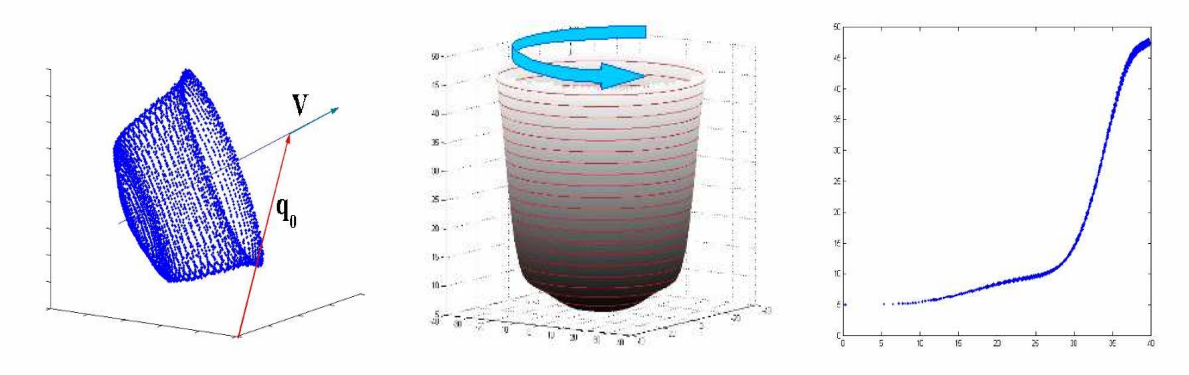

(a) (b) (c)

Figure 1: Axially Symmetric Geometry
(a) collection of range points (b) surface with iso-contours oriented about the axis of symmetry $(\vec{v}^t, \vec{q}_o^t)^t$ (c) 2D representation of the shape by means of its profile curve with respect to its axis as the z-axis.

3.1. 3D Surface Estimation Reduced to 2D Curve Estimation

The problem as stated requires the estimation of a 3D surface. However, as we have shown, a 3D axially symmetric surface is completely specified via its axis and corresponding 2D profile curve. In this section, we show that fitting a sequence of frustrums of cones to 3D data can be reduced to fitting a sequence of linear splines to 2D data. Fitting algebraic surfaces to 3D data also reduces to a 2D curve-fitting problem, but that reduction is more subtle and a bit more complex. We illustrate this by showing that the objective function to be minimized for 3D surface fitting reduces to an objective-function minimization for 2D data.

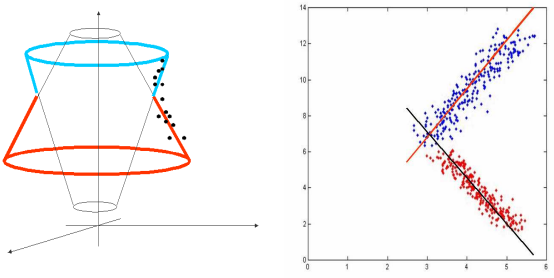

Figure 2: A model involving 2 line segments ($L = 2$)

1. Fitting a sequence of frustrums of cones. Consider Figs. 1(a) and 2 where there is a hypothesized axis specified by vector $(\vec{v}^t, \vec{q}_o^t)^t$, and in the latter figure, a hypothesized sequence of two frustrums of cones. The objective function to be minimized is the sum of squared *perpendicular distances* from the data points to the hypothesized surface. Hence, consider the i^{th} data point. This point lies in the plane determined by the point and the hypothesized surface-axis, and the perpendicular distance is measured within this plane. This perpendicular distance is simply the perpendicular distance to the hypothesized profile curve. Hence, consider a local orthogonal coordinate system (r, z) where the z line is the hypothesized surface axis and r is the distance to a point in the direction perpendicular to the z line. In this coordinate system, the i^{th} data point has coordinate vector $(r_i, z_i)^t$. Then 3D surface fitting becomes 2D linear spline fitting to a scatter of 2D data points in an (r, z) plane as in Fig. 2. In 2D, the error measure is the sum of squared distances from each data point to its appropriate straight-line segment.

2. Fitting an axially-symmetric algebraic surface. If the objective function used in fitting were simply algebraic distance, arguments analogous to those for fitting a sequence of conic frustrums, case 1 above, lead to the conclusion that the 3D objective function is equal to the sum of squared algebraic distances for algebraic-curve fitting to the scatter of points in a plane as in case 1. However, an additional complication is that for the i^{th} data point the Gradient-1 objective function (see section 3.3) includes a quadratic term involving the unit vector, $\vec{n_i}$, which is perpendicular to the measured 3D surface at the i^{th} data point. This unit normal is the unit normal to a 3D hyperplane fit to the measured data in a small neighborhood of the i^{th} data point. In general, this unit vector will not lie in the plane determined by the i^{th} data point and the hypothesized surface axis. Hence its contribution depends on its orthogonal projection into this plane and its component orthogonal to the plane. The equivalent objective function for 2D algebraic curve fitting involves the

$$E = \sum_{i=1}^{M} \sum_{l=1}^{N_i} \left| [z_{il} - z_i] \sin\theta_i - \left[r_{il} - \left(r_1 + \sum_{q=1}^{i} \tan(\theta_q)(z_{q+1} - z_l) \right) \right] \cos\theta_i \right|^2 \quad (4)$$

projection of the unit vector into the plane and this projection has length less than 1. The resulting objective function is slightly different than the objective function used in standard Gradient-1 algebraic curve fitting [8].

3.2. Fitting Local Spline Surface Models

With this model, we are estimating the profile curve as a continuous sequence of L piecewise linear segments. This is equivalent to estimating the 3D surface as a sequence of conic frustrums (see Fig. 2).

In this treatment, the 2D (r, z) data space is divided into L layers of equal size. Each layer is the subset of data points lying between a pair of lines perpendicular to the z axis, see Fig. 3.

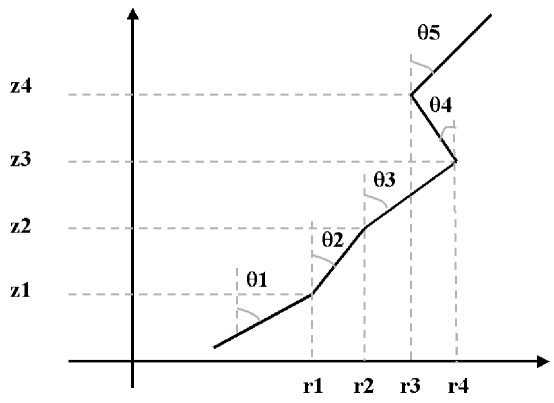

Figure 3: Continuous sequence of linear segments $(L = 5)$.

Note, a low-computation linear least squares fit to each data layer is possible by fitting lines independently to each layer, but the resulting lines may then not intersect at the layer boundaries, i.e. at $z_1, z_2, \ldots$. The resulting representation is not as good as that for which line segments are constrained to intersect at the layer boundary points z_i. We use the constrained fitting where the error minimized is the sum of squared distances from each data point to the line segment in its layer. The parameters to be estimated are one radius value, e.g., r_1, and the angles for each line segment, $\theta_1, \theta_2, \ldots, \theta_5$ in Fig. 3. The objective function to be minimized for M layers and N_i points in layer i is (4) (above) where (z_{il}, r_{il}) are z and r components respectively of the l^{th} data point. Unfortunately, this requires non-linear numerical minimization over the parameters $(r_1, \theta_1, \ldots, \theta_M)$.

Further improvement is obtained by estimating an optimal set of layer boundary points $z_1, \ldots, z_M$. In this case the strategy to be followed consist of a hierarchical two-step minimization procedure.

1. First keep a partitioning configuration $\{z_1, \ldots, z_M\}$ fixed, and find the optimal line segment sequence that fits the data according to such partitioning. This means to optimize with respect to $\{r_1, \theta_1, \ldots, \theta_M\}$.
2. Then correct the Z-partitioning in order to minimize the global error.

These two steps have to be iterated until convergence.

3.3. Global Algebraic Surface Models

Implicit polynomial curves provide a compact and low computational-cost method of representing shape [8, 9]. The explicit linear least-squares solution fitting methods that have been developed provide stable and robust polynomial fits in the presence of noise [8, 2]. We apply the Gradient-1 fitting algorithm recently developed by Tazdizen *et. al.* [8]. For a profile curve model of degree n there are $p = \frac{(n+1)(n+2)}{2}$ unknown coefficients. The form in (3) can be represented as the inner product of an unknown coefficient vector, Z, and a known monomial vector Y:

$$Z = \begin{bmatrix} a_{00} & a_{10} & a_{01} & \cdots & a_{n0} & a_{(n-1)1} & \cdots & a_{0n} \end{bmatrix}^t$$

$$Y = \begin{bmatrix} 1 & r & z & \cdots & r^n & r^{(n-1)}z & \cdots & z^n \end{bmatrix}^t$$

Using this representation, (3) is:

$$f_n(r,z) = Y^t Z$$

The Gradient-1 fitting algorithm makes use of the gradient of $f_n(r, z)$ in order to get a more stable estimate of the curve. As described in section 3.1 case 2, for the i^{th} 3D data point there is an associated estimated surface normal $\overrightarrow{n_i}$. This normal can be decomposed into a component n_i^p in the (r, z) plane and a component orthogonal to the plane. Note that $\|n_i^p\| \leq 1$. We use a modified version of the energy function (5) defined in [8] for estimating the profile curve coefficients.

$$e_{grad} = \sum_{i=1}^{m} (f_n(r_i, z_i)^2 + \mu \, \|n_i^p - \nabla f_n(r_i, z_i)\|^2) \quad (5)$$

where

$$\nabla f_n(r_i, z_i) = \nabla Y^t Z = \begin{bmatrix} \frac{\partial Y}{\partial r} \\ \frac{\partial Y}{\partial z} \end{bmatrix}_{2\times p} \begin{bmatrix} a_{00} & \cdots & a_{0n} \end{bmatrix}^t_{p\times 1}$$

In (5), μ is a weighting parameter which provides a more stable estimate $f_n(r,z)$ by forcing $f_n(r,z)$ to have a gradient close to the normal n_i^p. As shown in [8], this is a linear least squares problem which may be put in matrix form and has an explicit solution incurring little computational cost. In summary, fitting a 2D algebraic curve in this reference plane is equivalent to specifying the 3D surface.

4. Estimation Computation-Algorithm

The non-linear iterative minimization necessary for estimating the surface parameters is similar for both methods 3.2 and 3.3. Each iteration may be decomposed into 2 steps:

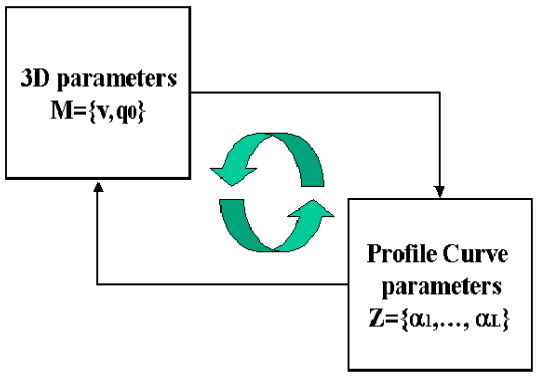

Figure 4: Two-step iterative minimization procedure.

1. Based on the value of the objective function after the preceding iteration, choose a new-value for the parameter vector $(\vec{v}, \vec{q}_o)$ specifying the axis of symmetry.

2. Estimate the best parameter values for the profile curve with respect to this axis using the method described in (3.2) or (3.3).

The objective functions defined in both methods are highly non-linear. Consequently, convergence to a local minimum may occur if minimization is started far from the true parameter value. The initialization of the estimation algorithm needs only a hypothesized axis of symmetry in order to begin. We use the initialization given in the next section.

4.1. Initialization

The *initial estimate* of the axis of symmetry is computed as an *explicit expression* resulting from a linear least squares formulation in Plücker coordinates as given by Pottmann *et. al.* [7]. In this approach, the authors consider the set of normal lines each defined by a data point and its corresponding normal (see Fig. 5). If each point lies on the axially symmetric surface, each of these lines will intersect the axis of symmetry at some point. Unfortunately, the lines do not intersect the true axis either because the measurement data is noisy and does not lie on the true surface or because the surface is not exactly axially symmetric. The goal of the authors is to find the straight line that has minimum distance to this set of normal lines. Using Plücker coordinates, the authors present a linear least-squares solution to this problem involving SVD of a covariance matrix.

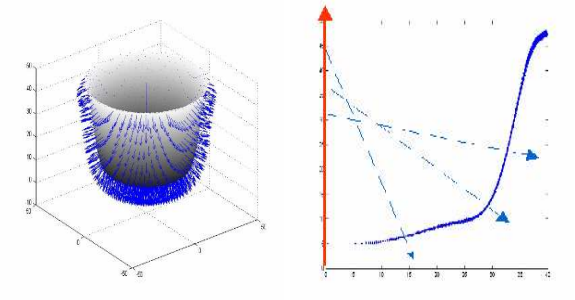

Figure 5: Plücker coordinate normal line solution

Unfortunately, the accuracy of this method depends entirely upon the reliability of the estimated normal vectors. Results may be poor in the presence of noise, if normal lines exist which are approximately parallel to the axis of symmetry, or if all normal lines roughly intersect at a point, as happens if the measured surface patch is on a sphere. For this reason, we use this solution only as a low computation-cost initial starting point. Hereafter, the algorithms applied make use of both the measured range points and the surface normals to refine the initial guess.

5. Experimental Results

Experiments were performed using data obtained from a 3D CT scan of a pottery sherd (i.e. fragment). The outer surface points were isolated and used as input to the estimation algorithm using method 3.2 with 5 linear spline segments (see Fig. 6).

As shown in Fig. 7(a), the linear segments effectively represent the the local shape of the data. Adaptive subdivision of the space allows us to exploit any local linearities in the surface which can produce very efficient representations when dealing with objects which may include flat

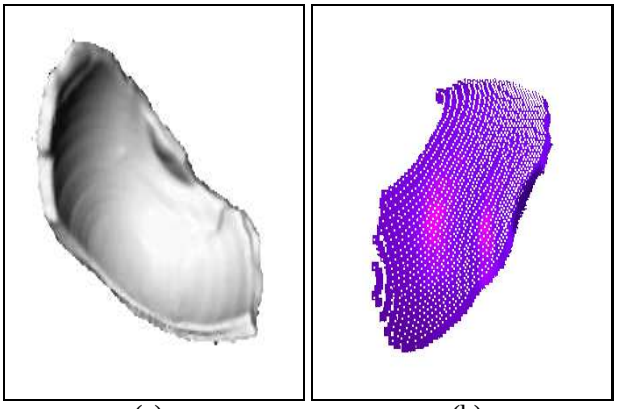

(a) (b)

Figure 6: (a) image of CT scan data (b) isolated outer surface data (shown in white).

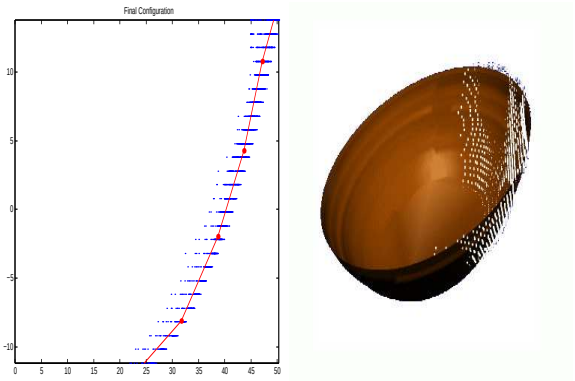

(a) (b)

Figure 7: (a) estimated linear spline profile curve using method of section 3.2 ($L = 5$) (b) axially symmetric solution via method 3.2 with 3D data overlaid (in white) on the solution.

areas. Ideally, the data scatter about the estmated profile curve would have 0 extent. The fact that it doesn't is because there is some noise in the measured data or the surface is not perfectly axially symmetric or the axis estimate is slightly in error. Fig. 7(b) shows the surface associated with the estimated linear spline model transformed into the original coordinate system of the data. Of particular note is the fact that the bump on the upper rim of the object did not significantly effect the surface estimate, demonstrating some robustness to surfaces which are not perfectly symmetric.

Figures 8 and 9 illustrate that if the data patch is small, one may have to consider two or more pairs of axes and corresponding profile curves for the data patch, and choose the appropriate pair later using additional information. We compute these multiple solutions by using, when appropriate, up to two suboptimal initial solutions obtained from the

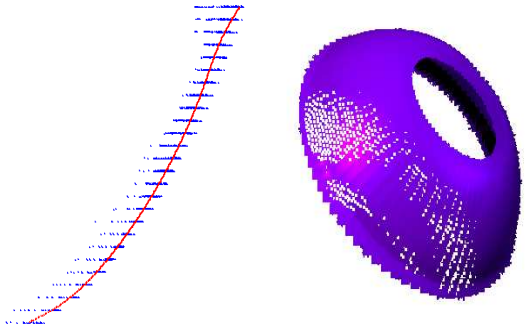

(a) (b)

Figure 8: Solution 1 (a) estimated 5^{th} degree profile curve using method of section 3.3 (b) axially symmetric solution via method 3.3 with 3D data overlaid (in white) on the solution.

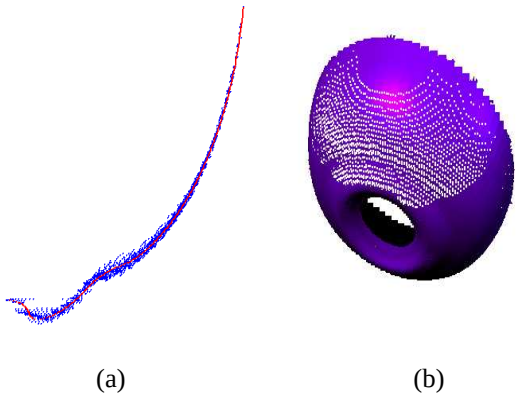

(a) (b)

Figure 9: Solution 2 (a) estimated 5^{th} degree profile curve using method of section 3.3 (b) axially symmetric solution via method 3.3 with 3D data overlaid (in white) on the solution.

approach in section 4.1.

A second experiment was performed using data obtained from Vitana's ShapeGrabber laser range finder (see Fig. 10). As shown in Fig. 10(a), the subject was an object with rough surface texture and imperfect symmetry. Before scanning, all of the object save a small portion of the upper half was occluded from the measuring device as shown in Fig. 10(b). The resulting data represented both the occluding objects and our test object. The data pertaining to our test object was segmented such that only surface points of interest remained (see Fig. 11(a)) which were then used as input to the estimation algorithm using method 3.3. As shown in Fig. 11(b), the resulting surface reconstruction fits

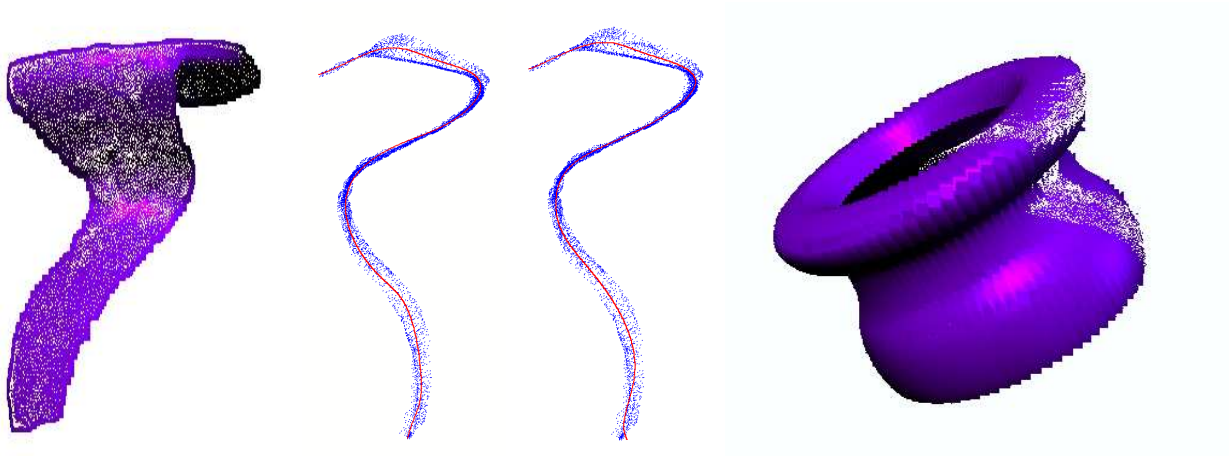

(a) (b) (c) (d)

Figure 11: (a) isolated patch of range data (in white) (b) estimated 5^{th} degree profile curve (c) estimated 6^{th} degree profile curve (d) reconstructed 3D surface with data superimposed (also in white).

(a) (b)

Figure 10: (a) experimental object (b) object is occluded, and only a small surface patch is observable to the 3D scanner.

closely the original range data. Evident in model is that the estimated curve closely models the local curvature of the surface. In addition, this model has the benefit of being a much more compact representation of the shape. The 5^{th} degree 2D IP requires only 21 parameters to represent the global shape of the object of interest which clearly has some complex surface geometry. Note that $r(z)$, the radius function, is multivalued because the top of the pot and immediate inside surface are modelled. The estimated geometry is completely captured by an estimated axis and associated profile curve for the data patch.

6. Summary and Conclusions

We have introduced two new methods for simultaneously estimating the pose and shape of a small patch from an axially symmetric object using 3D data that does not include the self occluding boundary. These methods generate accurate and compact representations of a possibly complex and imperfect object. For illustrative purposes, experiments have been presented with each method using different kinds of input data (i.e. 3D CT scan data and 3D laser range data), demonstrating that the methods are practical and accurate, even when the data patch subtends only a small angle around the object axis. In addition, although some nonlinear minimization is necessary, the computational cost of these methods are not excessive. An extension in progress of this work is computation of a covariance matrix for the estimated patch axis and profile curve. This can be obtained approximately from the Hessian of the objective function minimized in the estimation. We see this work as a step along the very important path of pulling out useful geometric structure from clouds of unorganized 3D data points.

References

[1] http://www.shapegrabber.com. ShapeGrabber Inc.

[2] M. Blane, Z. Lei, H. Civi, and D. Cooper. The 3L algorithm for fitting implicit polynomial curves and surfaces to data. *PAMI*, March 2000.

[3] G. Medioni and Z. Mourad. *Object Representation in Computer Vision*, chapter The Challenge of Generic Object Recognition, pages 217–232. Springer, 1994.

[4] T. Nakamura, A. M., and S. Yoshiaki. A qualitative approach to quantitative recovery of SHGC's shape and pose

from shading and contour. In *Proc. IEEE Conf. on Comp. Vision and Pattern Recognition*, pages 116–122, 1993.

[5] R. Nevatia and T. Binford. Description and recognition of curved objects. In *Artificial Intelligence*, pages 77–98, 1977.

[6] R. Nevatia and F. Ulupinar. Shape from contour: Straight homogeneous generalized cylinders and constant cross section generalized cylinders. *PAMI*, pages 120–135, February 1995.

[7] H. Pottmann, M. Peternell, and B. Ravani. An introduction to line geometry with applications. *Computer-Aided Design*, 1999.

[8] T. Tasdizen, J. Tarel, and D. Cooper. Improving the stability of algebraic curves for applications. *IEEE Trans. on Image Proc.*, pages 405–416, March 2000.

[9] G. Taubin. Estimation of planar curves, surfaces and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation. *PAMI*, 1991.

[10] L. VanGool, M. Pollefeys, and M. Proesmans. Self calibration and metric reconstruction in spite of varying and unknown internal camera parameters. In *Proc. International Conference on Computer Vision*, pages 90–95, 1998.