Alignment of Multiple Non-overlapping Axially Symmetric 3D Datasets

Andrew Willis, David B. Cooper
Brown University, Division of Engineering, Providence, RI 02912, USA

Abstract

Uknown to us, an axially-symmetric surface is broken into disjoint pieces along a set of break-curves, i.e., the curves along which the surface locally breaks into two pieces. A subset of the pieces are available and for each of them we obtain noisy 3D measurements of its surface and break-curves. Using the piece measurements and knowledge of which pieces share a common break-curve, we propose a stochastic method for automatically estimating the unknown axially-symmetric global surface. Surface and break-curve estimation is then an alignment problem where we must estimate the unknown axially-symmetric surface and break-curves while simultaneously estimating the Euclidean transformation that positions each measured piece with respect to the a-priori unknown surface. Parameter estimation is implemented as maximum likelihood estimation where we seek the global pot geometry which best explains the measured fragment data. This new approach is robust, fast, and accurate. Experimental results are presented which solves an application of interest, specifically the reconstruction of archaeological pots from subsets of their surface pieces.

1 Problem

We address the problem of estimating an unknown free-form axially symmetric surface from an incomplete set of its pieces. The pieces individually describe small portions of the unknown surface and are connected to each other by break-segments, i.e., locations where the pieces were broken apart. A surface reconstruction program must determine which break-segments on individual pieces should be connected in order to estimate the surface. This is a difficult search problem whose solution is given in [7, 8]. We assume that the correspondence between the piece break-segments is already known and concentrate on the problem of simultaneously estimating the geometric parameters of the unknown axially symmetric surface while aligning the pieces.

2 Parameters to Estimate

The break-curve parameters, $\boldsymbol{\beta}$, are points on the surface where the surface has broken into two pieces. These locations include *vertices*—locations of Y and T junctions (see Fig. (1)). Note that junctions denote points which are high-curvature points on piece boundaries. Consider the piece boundaries in the vicinity of a junction. T-junctions are points which are high-curvature points on two of the three piece boundaries and Y-junctions denote points which are high-curvature points for all three piece boundaries. We also refer to these high-curvature points on piece boundaries as *vertices.* The points that constitute our representation for the break curves, referred to as a *break-point segments,* is a sequence of K points for each curve starting with the vertex. The points in each sequence occur at successive intervals of fixed length from the vertex (see Fig. (2)). These points along with a surface normal at each point constitute $\boldsymbol{\beta}$—our parameterization of the surface-curves for the axially-symmetric surface. Hence, break-point segment v is written $\boldsymbol{\beta}_v = ((\mathbf{p}_1, \mathbf{p}_2, \ldots \mathbf{p}_K), (\mathbf{n}_1, \mathbf{n}_2, \ldots, \mathbf{n}_K))$ where $\mathbf{p}_k$ denotes the k^{th} 3D point and $\mathbf{n}_k$ denotes the k^{th} 3D normal for $\boldsymbol{\beta}_v$. The group of all break-curve parameters is $\boldsymbol{\beta} = \cup_{v=1} \boldsymbol{\beta}_v$.

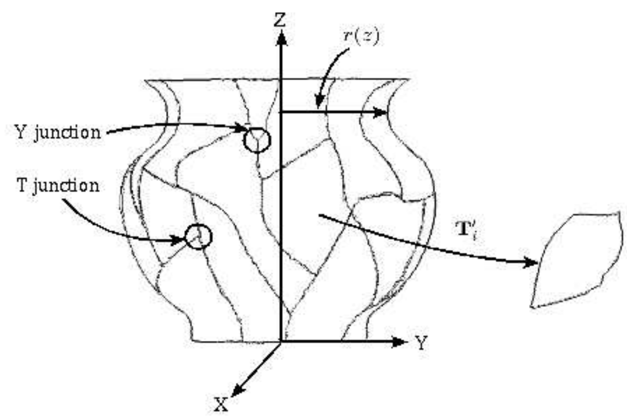

Figure 1: A broken vessel and its parametric description.

The k^{th} point in a break-point segment is the location where the break-curve intersects a sphere of radius kR centered at a vertex point where $k = [0, 1, 2, \ldots, K-1]$ and R is the radius of a sphere centered at the vertex. When we measure break-point segment data, noisy estimates of these points are extracted starting with a hypothesized vertex point (see §3). Figure (2) illustrates two break-point

from the Proceedings of ICPR (International Conference on Pattern Recognition) Vol IV, August 2004

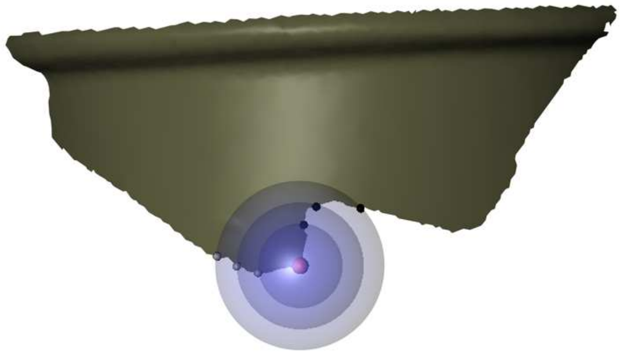

Figure 2: Break-point segments : a piece data-set, in grey, and one of the piece vertices is shown as a large opaque red sphere. Two sets of break-point data which we call *break-point segments* are generated shown as light grey and black points on the break-curve. Each break-point segment shown has 4 ordered elements, starting with the vertex and then listed in order of increasing distance from the vertex. .

segments for a vertex where $K = 4$.

We have parameterized the surface axis of symmetry using the standard parametric equation of a 3D line as shown in (1).

$$\begin{array}{l} x = m_x z + b_x \\ y = m_y z + b_y \end{array} \quad (1)$$

Hence $\mathbf{l} = (b_x, b_y, m_x, m_y)$ and consists of the pair (m_x, m_y), specifying the slope of the line when it is projected onto the xz-plane and the yz-plane, respectively, and the pair (b_x, b_y), specifying where the line intercepts the xy-plane at $z = 0$.

The profile curve $\alpha(r, z)$, with respect to the piece axis $\mathbf{l}$ (see Fig. 2), defines a *3D axially symmetric algebraic surface* with axis $\mathbf{l}$ where r denotes *radius*, i.e., the shortest distance between a 3D point and the axis and z denotes *height* along the axis. The surface parameters are the coefficients of the algebraic profile curve in (2) and the axis of symmetry in (1). The form of the 2D algebraic profile curve of degree d is (2).

$$\alpha(r, z) \quad = \sum_{0 \leq j+k \leq d;\ j,k \geq 0} \alpha_{jk} r^j z^k = 0 \qquad (2)$$

Hence $\boldsymbol{\alpha} = (\cup_{j,k} \alpha_{jk})$ is the vector of coefficients for the implicit polynomial curve of degree d. Note that for our experiments, $d = 6$.

We assume that each piece undergoes an arbitrary rigid Euclidean transformation which moves the piece to its measurement position. Hence, for the i^{th} piece we must estimate the transformation, $\mathbf{T}_i$, which moves the piece from its measurement position to its aligned position. We parametrize a 3D transformation with 6 parameters consisting of 2 parts : (1) a 3D translation vector $\mathbf{t}$, and (2) a 3D rotation $\mathbf{R}$. The 3x3 rigid rotation matrix $\mathbf{R}$ is represented using the so called axis-angle parameters which describe rotation in terms of a rotation angle ψ about a 3D unit vector $\mathbf{n_R}$. Hence our rotation is the 3D vector $(\psi \mathbf{n_R})$ and we refer to the equivalent 3x3 rotation matrix as $\mathbf{R}$. The 3D transformation parameters are $\mathbf{T} = (\mathbf{t}, \psi \mathbf{n_R})$ (see Appendix A of [4] for additional details on this parameterization). Further, we denote transformations of piece datasets as $\mathbf{T}(\mathcal{D})$ for surface point/normal data and $\mathbf{T}(\mathcal{B})$ for break segment point/normal data. In this notation it is assumed that $\mathbf{T}(\mathcal{D})$ indicates that a single transformation, $\mathbf{T}$, operates on each of the points and normals in $\mathcal{D}$ according to (3).

$$\begin{array}{rcl} \mathbf{Tp} & = & \mathbf{Rp} + \mathbf{t} \\ \mathbf{Tn} & = & \mathbf{Rn} \end{array} \quad (3)$$

3 Data Generation Model

The surface measurement data is provided by a Shapegrabber laser/camera scanner [1]. It produces 15,000 3D points/sec. at a resolution and accuracy of the order of 0.25mm. The data is divided into two sets : (1) 3D surface points and normal measurements,, denoted $\mathcal{D}_i$ for the i^{th}surface piece, and (2) 3D break-segment points and their surface normals, denoted $\mathcal{B}_v$ for the v^{th} break-segment.

Data for each break-segment is extracted according to a simple parameterization of the unknown surface break-curves adopted in §2. Hence, break-point segment data is simply a special sequence of K measured 3D points and outer surface normals (see the discussion on break-curve parameters in §2 and Fig. (2) for clarification).

Assumptions

Surface measurement points are *i.i.d.* $\mathbf{N}(0, \sigma_{\mathcal{D}}^2 \mathbf{I})$
These are independent, identically distributed, spherically symmetric Gaussian perturbations in 3-space about each point on the true surface with mean 0 *and variance* $\sigma_{\mathcal{D}}^2$. See [2] for a justification of this model.

Surface measurement normals are *i.i.d.* $\mathbf{N}(0, \lambda \sigma_{\mathcal{D}}^2 \mathbf{I})$
These are independent, identically distributed symmetric Gaussian perturbations on the unit sphere , i.e., in SO(3), about the true surface normal for each point on the sherd outer surface with mean 0 *and variance* $\lambda \sigma_{\mathcal{D}}^2$. These are independent and for the j^{th} piece are distributed over a spherical cap about a mean that is normal to the surface as represented by an axis/profile-curve for the j^{th} piece.

Break-segment measurement points are *i.i.d.* $\mathbf{N}(0, \sigma_{\mathcal{B}}^2 \mathbf{I})$
These are independent, identically distributed spherically symmetric Gaussian perturbations in 3-space about each point on the true break-curve, with mean 0 *and variance* $\sigma_{\mathcal{B}}^2$. Note that more appropriate but more complicated models can be used.

4 Algorithm

Alignment is done by adding one piece at a time, each addition producing a new *configuration*. For each new configuration, *all* parameters are re-estimated to minimize the joint probability density for all the break-point segment and surface data associated with the pieces in the configuration. However, this is not done in one simultaneous nonlinear minimization, since that would be both much too time consuming and would converge to a local maximum. Note, for an N piece configuration, the number of parameters to be estimated is : $6(N-1)$ Euclidean transformation parameters; $4+\frac{d(d+1)}{2}$ global surface parameters for a d^{th} degree surface (if $d=6$ then there are 29 parameters); and a minimum (and likely many more) of $3K(N-1)$ break-point segment parameters where K is the number of break-points in a segment. Estimating an axis/profile curve for a surface data set of a few thousand points may require up to 1 minute if using an arbitrary starting point. Using a small-error starting point can result in orders of magnitude computation reduction. Estimating a Euclidean transformation for aligning a pair of matched break-point segments is a linear least squares computation on the order of a millisecond.

1. Given a N piece configuration denoted $ConfigN$, we begin by treating the configuration as though a single piece. We estimate the global axis/profile-curve parameters for $ConfigN$ by solving the non-linear equation (4) (see §4.1 for details).

$$\widetilde{\mathbf{l}}, \widetilde{\boldsymbol{\alpha}} = \arg\max_{\mathbf{l},\boldsymbol{\alpha}} \ln\left(\prod_{\mathcal{D}_i \in ConfigN} \mathbf{P}\left(\mathcal{D}_i | \mathbf{l}, \boldsymbol{\alpha}, \widetilde{\boldsymbol{\beta}}, \widetilde{\mathbf{T}}_i\right)\right) \tag{4}$$

   Upon convergence, we store the parameters of the estimated axially symmetric surface, $(\widetilde{\mathbf{l}}, \widetilde{\boldsymbol{\alpha}})$.

2. We take a new piece $N+1$ and find the Euclidean transformation that moves this piece so that its break-point segments align with the matched break-point segments of $ConfigN$ and its surface data fits the 3D surface model for $ConfigN$ by maximizing (5) (see §4.2).

$$\widetilde{\mathbf{T}}_{N+1}, \widetilde{\boldsymbol{\beta}} = \arg\max_{\mathbf{T}_{N+1},\boldsymbol{\beta}} \ln\left(\mathbf{P}\left(\mathcal{D}_{N+1}, \mathcal{B}_{N+1}, \mathcal{B}_{cN} | \widetilde{\mathbf{l}}, \widetilde{\boldsymbol{\alpha}}, \boldsymbol{\beta}, \mathbf{T}_{N+1}\right)\right) \tag{5}$$

   Where $\mathbf{T}_{N+1}$ denotes the Euclidean transformation for piece $N+1$, $\mathcal{D}_{N+1}$ denotes transformed surface data of piece $N+1$. $\mathcal{B}_{cN}$ and $\mathcal{B}_{N+1}$ denote the matched break-point data segments of $ConfigN$ and the transformed piece $N+1$ respectively. $\boldsymbol{\beta}$ denotes the unknown parameters of the matched break-point segments.

3. Add piece $N+1$ to $ConfigN$ and go to step (1) if the number of pieces in $ConfigN$ is smaller than the total number of available pieces.

It is possible to jiggle all N pieces a bit by re-estimating the transformation for piece 1 to make its data better fit the configuration defined by piece 2 through N, then repeat with piece 2, etc., but the improvement is small, and doesn't appear to be necessary.

4.1 Estimating the Surface

Surface shape parameters are obtained by finding the axially symmetric algebraic surface which best fits the aligned data sets using the method outlined in [9]. This is a non-linear minimization consisting of two steps:

1. Based on the value of the objective function after the preceding iteration, choose a new value for the axis parameters, $\widetilde{\mathbf{l}}$.

2. Fixing the axis $\mathbf{l} = \widetilde{\mathbf{l}}$, we then find the axially symmetric surface $\boldsymbol{\alpha}$ which maximizes the probability of the data by fitting an axially symetric surface to the 3D data with axis $\widetilde{\mathbf{l}}$. This is equivalent to solving the weighted linear least-squares problem (6) which has an explicit solution and incurs little computational cost.

$$e(\boldsymbol{\alpha}|\mathbf{l} = \widetilde{\mathbf{l}}) = \sum_{\mathbf{p},\mathbf{n} \in ConfigN} \left(\boldsymbol{\alpha}^2(\mathbf{p}) + \frac{1}{\lambda}\left\|\mathbf{n} - \nabla\boldsymbol{\alpha}(\mathbf{p})\right\|^2\right) \tag{6}$$

Since the surface model depends upon the axis, the resulting objective function is non-linear and is minimized directly using simulated annealing. Note that the error metric $\boldsymbol{\alpha}^2(\mathbf{p})$ is the so-called algebraic distance which can be used as an approximation of the Euclidean distance (for details see [5]), and $\frac{1}{\lambda}\|\mathbf{n} - \nabla\boldsymbol{\alpha}(\mathbf{p})\|^2$ is for regularization [9].

4.2 Estimating the Transformation

We obtain a coarse estimate for the piece $N+1$ transformation by aligning the subset of its corresponding break point segments to those of $ConfigN$. Given our noise assumptions from §3, this corresponds to solving (7).

$$\check{\mathbf{T}}_{N+1} = \min_{\mathbf{T}_{N+1}} \left\|\mathcal{B}_{cN} - \mathbf{T}_{N+1}(\mathcal{B}_{N+1})\right\|^2 \tag{7}$$

This is commonly referred to as the absolute orientation problem to which there exist several explicit linear solutions [6, 3]. We use the solution proposed by [6]. Note that (7) includes a term involving the surface normals at each break segment data point which aligns the tangent planes of corresponding points which improves the alignment significantly when the matched break-point segment data has little curvature.

Using surface estimate $(\widetilde{\mathbf{l}}, \widetilde{\boldsymbol{\alpha}})$ from step (1) of the algorithm, we solve the non-linear alignment problem (8) using our rough estimate $\check{\mathbf{T}}_{N+1}$, and $\check{\boldsymbol{\beta}} = \frac{\check{\mathbf{T}}_{N+1}(\mathcal{B}_{N+1}) + \mathcal{B}_{cN}}{2}$ as an initial point.

$$e(\mathbf{T}_{N+1}, \boldsymbol{\beta}) = \boldsymbol{\alpha}^2(\mathbf{T}_{N+1}(\widetilde{\mathcal{D}}_{N+1})) + \|\mathbf{T}_{N+1}(\mathcal{B}_{N+1}) - \boldsymbol{\beta}\|^2 + \|\mathcal{B}_{cN} - \boldsymbol{\beta}\|^2 \tag{8}$$

Where $\widetilde{\mathcal{D}}_{N+1}$ denotes the subset of the piece $N+1$ surface data as shown in Fig. (3b). The term $\boldsymbol{\alpha}^2(\mathbf{T}_{N+1}(\widetilde{\mathcal{D}}_{N+1}))$ is the approximate Euclidean distance of the transformed data set $\widetilde{\mathcal{D}}_{N+1}$ to the estimated axially symmetric surface

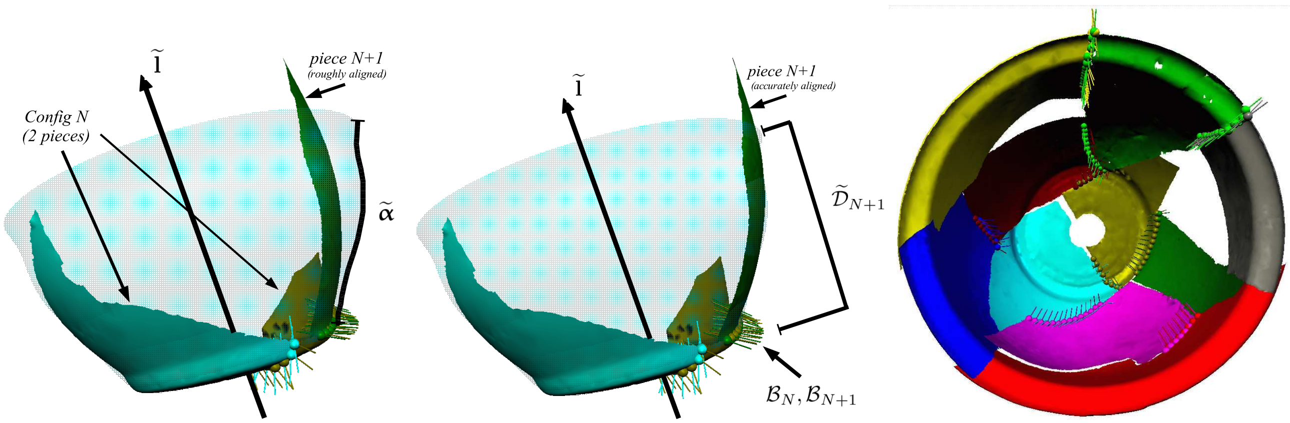

(a) Example of algorithm Step 1. (b) Example of algorithm Step 2. (c) A 10 piece alignment.

Figure 3: Surface Alignment : $ConfigN$ consists of two pieces as shown in (a). The axially symmetric surface $(\widetilde{\boldsymbol{\alpha}}, \widetilde{\mathbf{l}})$ obtained from step (1) of our algorithm is shown in transparent cyan and the profile $\widetilde{\boldsymbol{\alpha}}$ is shown as a black curve. Piece $N+1$, in green, is added to the configuration using the rough alignment given by solving (7), the resulting 3-piece configuration is shown in (a). We then solve (8) to obtain an accurate alignment for piece $N+1$ using its matched break point data segments, $\mathcal{B}_{N+1}$, and a subset of its surface points, $\widetilde{\mathcal{D}}_{N+1}$. $\widetilde{\mathcal{D}}_{N+1}$ denotes those surface points of piece $N+1$ whose height along the axis $\widetilde{\mathbf{l}}$ is within the height interval defined by the Config N data set, i.e., the bracketed region in (b).

for $ConfigN$. The surface alignment error being minimized here is the square of the algebraic distance (see Section XI of [5] for details where a similar alignment problem is solved). The break-segment alignment error is $\|\mathbf{T}_{N+1}(\mathcal{B}_{N+1}) - \boldsymbol{\beta}\|^2 + \|\mathcal{B}_{cN} - \boldsymbol{\beta}\|^2$ which denotes the Euclidean distances between the matched breakpoints on piece $N+1$ and $ConfigN$ and the true locations of the break points $\boldsymbol{\beta}$. The solution to (8) is computed by direct non-linear minimization using the well-known Levenberg-Marquardt algorithm.

5 Results

Given the large amount of unknown parameters being estimated, our algorithm succeeds in finding the global MLE solution quickly. The 3 piece configuration in Fig. (3b) required approximately 10 seconds to compute. Note that this minimization involves 221 parameters (12 transformation parameters, 29 surface parameters, and 180 break segment parameters). Fig. (3c) shows the solution for a 13 piece axially symmetric vessel where only 10 pieces are available which required approximately 45 seconds to compute, the number of estimated parameters here is 893 (54 transformation parameters, 29 surface parameters, and 810 break segment parameters).

6 Conclusions

We have explained and demonstrated a robust and fast method for estimating an unknown axially symmetric surface from measurements of a subset of its pieces. As others have in the past, we decompose the problem into a set of $N-1$ pairwise piece alignments. Our method is unique in that it solves the problem using a 2 step recursive algorithm. In step 1, the geometric constraint of axial symmetry is used to obtain an estimate of the unknown surface for a configuration of aligned pieces. Step 2 uses this surface estimate to obtain accurate estimates of the break segment and transformation parameters for a new piece added to the configuration. The algorithm is shown to work and performs well even though dealing with a vast number of unknown parameters. This material is based upon work supported by the National Science Foundation under Grant No. 0205477.

References

[1] Shapegrabber inc. http://www.shapegrabber.com.

[2] R. Bolle and D. B. Cooper. On optimally combining pieces of information, with application to estimating 3-D complex-object position from range data. *PAMI*, 8(5):619–638, 1986.

[3] B. Horn. Closed-form solution of absolute orientation using unit quaternions. *Journal Optical Society of America*, 4(4):629–642, 1987.

[4] X. Pennec and J. Thirion. A framework for uncertainty and validation of 3-d registration methods based on points and frames. *IJCV*, 1997.

[5] G. Taubin. Estimation of Planar Curves, Surfaces, and Nonplanar Space Curves Defined by Implicit Equations with Applications to Edge and Range Image Segmentation. *PAMI*, 13(11):1115–1138, 1991.

[6] S. Umeyama. Least-Squares Estimation of Transformation Parameters Between Two Point Patterns. *PAMI*, 13(4):376–380, 1991.

[7] A. Willis, D. Cooper, et al. Assembling virtual pots from 3D measurements of their fragments. In *VAST*, pages pp. 241–253, 2001.

[8] A. Willis, D. Cooper, et al. Bayesian Pot-Assembly from Fragments as Problems in Perceptual-Grouping and Geometric-Learning. In *ICPR*, volume III, pages 297–302, 2002.

[9] A. Willis, D. Cooper, et al. Accurately Estimating Sherd 3D Surface Geometry with Application to Pot Reconstruction. In *CVPR Workshop : ACVA*, June 2003.