

Estimating a-priori Unknown 3D Axially Symmetric Surfaces from Noisy Measurements of their Fragments

Andrew Willis
Dept. of Electrical and Computer Engineering
University of North Carolina
Charlotte, NC 28223

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

Abstract

In this paper, we present a computationally efficient technique for solving the difficult problem of estimating the global shape of a ceramic pot from measurements of its fragments. Each unknown pot is modeled as a surface of revolution, i.e., a 3D line—the central axis of the pot—and a 2D profile curve with respect to that axis. For each fragment, a probabilistic distribution is estimated which models both the geometric shape of the fragment and the variability of the estimated fragment shape. Estimation of the global pot shape is then a Maximum Likelihood Estimation (MLE) problem where we seek the values of the Euclidean transformation parameters that maximize the joint probability of the matched fragments' axis/profile-curve models (which includes the additional constraint that the matched fragments must share a common central axis). This is a new type of curve-analysis problem and our solution is a new and effective approach applicable for generic constrained 2D curve alignment and for modeling of 3D axially-symmetric surfaces and for comparing geometric models which may correspond over a subset of the complete model.

1 Introduction

Automatic assembly of 2D puzzles has been a topic of interest for many years with seminal work in the area represented by [1] in which there is an effort to automatically assemble jigsaw puzzles from their pieces. Yet, the problem of assembling *real world* puzzles, i.e., puzzles generated by nature (not man-made) has received little attention until recently. Here, fragments are irregular and can match along any subset of their complete boundary. Additionally, typical real world puzzles introduce a vast variety of possible sources of confusion. Some of these include (1) physical degradation of the fragments due to chipping and, if exposed to the elements, erosion, (2) the number of puzzles being solved may not be known, i.e., we are given a collection of fragments coming from an unknown number of

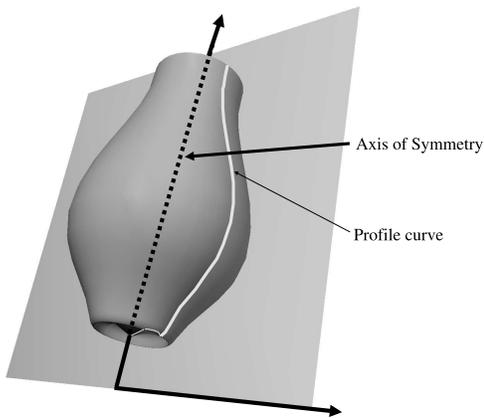


Figure 1: Surface of Axial Symmetry

objects, (3) fragments may be missing as they may have yet to be discovered or may have been destroyed by some physical phenomenon. Progress on this problem is important in fields such as forensics, archaeology, and anthropology where the investigator may not know the original shape of the broken object or the specific positions of the fragments in the original shape. Good progress on real world 2D puzzles is shown in works such as [2, 3]. Initial work on automatically assembling 3D puzzles is given in [4, 5]. Yet, automatic assembly of real world 3D puzzles is a difficult and computationally intensive process [6]. This is due to the complexity of matching 3D surfaces and compounded by the very large number of fragment comparisons needed in even modest 3D puzzles. The computational killer here is that published approaches [7] process the raw range data each time a merge of two or more surface fragments is hypothesized and estimated, and the number of data points for a fragment is typically in the many thousands.

In this paper, we propose a new technique for reconstructing the profile curve (see fig. 1) of an archaeological pot given measurements of its fragments. Fragment data is generated using a dense-data 3D laser range scanner that

generates range images of the fragments by scanning their outer surface. Range images are typically 3D triangular surface meshes, i.e., a connected group of (x, y, z) surface measurements where the number of measurements depends on the size of the fragment (typically 10k-200k points).

The approach uses a stochastic geometric model for each fragment's, 3D surface in terms of an axis/profile-curve pair. Note that this model consists of parameters to describe a 3D line, the axis of symmetry, and parameters to describe a 2D curve with respect to the axis and is *completely equivalent* to modeling a 3D surface of axial symmetry. Each axis/profile-curve pair characterizes some portion of the complete unknown global axis/profile-curve, in terms of a curve-fragment and a preferred coordinate system for the curve fragment, i.e., the coordinate system defined by the estimated 3D axis of symmetry for the fragment. We then combine axis/profile-curve models by matching similar portions of multiple curve fragments by computing the maximum likelihood estimate of the complete 3D surface from a set fragment surface models.

The approach differs from past approaches in two important ways :

1. Our approach matches fragments together based only on their estimated parametric distributions which is a distribution on a set of 2D points, the profile-curve, and their preferred coordinate system, defined by the axis of symmetry. Previous approaches combine fragments based on the data measured for each of the fragments. Since data for a single fragment includes up to 200k points, computation of the best fragment alignment is computationally costly.
2. Our approach makes no use of fragment boundary matching. Previous approaches seek to compute the shape of the vessel and simultaneously estimate the original positions of each fragment in the complete pot. Here, we seek only to extract the global unknown shape of the pot by aligning surface estimates obtained from each fragment.

A new stochastic curve model is introduced as an approach to the difficult problem of quickly matching and aligning 3D surface fragments. The stochastic model characterizes both the 2D geometric shape and shape variability of each fragment's axis/profile-curve using a probability distribution defined on a set of geometric parameters that summarize the fragment surface shape which typically has 20-50 parameters depending upon the fragment size. It is these stochastic axis/profile-curve models for the individual fragments rather than the raw data for individual fragments that we use in estimating the pot global axis/profile-curve. Since there are far fewer parameters than measured data points, the estimated models may be aligned much faster than techniques that use the measured data points with little loss of

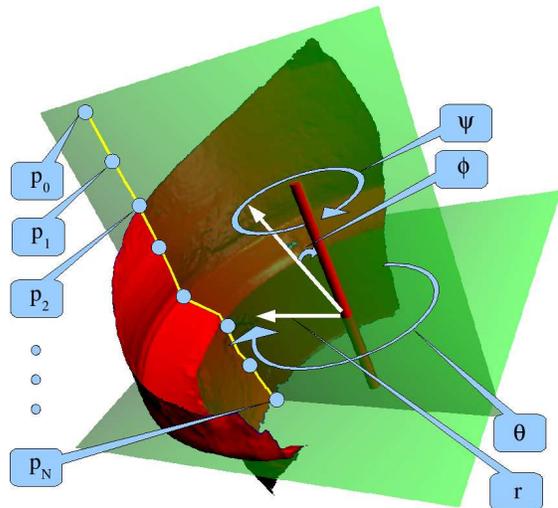


Figure 2: Parameters for the shape model (see §2 for details).

accuracy. We emphasize that the new problem should be viewed as one where we are given an unorganized pile of stochastic curve fragments, each with respect to an axis, and our job is to concatenate them, such that they may overlap, and such that the resulting curve (or curves) has the highest probability.

2 Fragment Model Parameters

For each fragment, a stochastic model is estimated from the sampled (x, y, z) surface data. The fragment shape model is an estimate of an axis of symmetry and a portion of some pots' unknown global profile curve, jointly. As previously mentioned, this geometry is specified in terms of a 3D line, the axis of rotational symmetry for the pot, and a 2D profile curve, the shape of the pot measured with respect to this axis. Consider an orthogonal (x, y, z) coordinate system. For convenience, our model is developed within a cylindrical coordinate system where each measured surface point has (r, θ, z) coordinates. e_z is a unit vector from the origin in the direction of the world coordinate system z -axis. e_r is a unit vector perpendicular to e_z lying in the x - y plane. With θ denoting an angle with respect to the x -axis and in the x - y plane any point in (x, y, z) can be specified in (r, θ, z) coordinates. In this coordinate system, a 3D axially symmetric surface may be completely specified by taking it's axis to be the z -axis and then specifying a 2D curve in the $(r, 0, z)$ plane, that is, the (r, z) plane with $\theta = 0$. Typically this curve is specified as a function, $r = f(z)$, along the z -axis. However, the shape of base and decorative rim fragments are often multi-valued for a given

z value making this parameterization inappropriate for our purposes. Hence, we use the commonly adopted arc-length parameterization which models the 2D curve as an ordered set of 2D points sampled at constant intervals of arc length along the curve. The fragment profile curve with N points is then denoted $\mathbf{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$ where the i^{th} point is a 2D point with coordinates $\mathbf{p}_i = (r_i, z_i)$.

3 Estimating Fragment Models

The axially symmetric surface from which each measured fragment is generated is *a-priori* unknown. Hence, the parameters Θ must be estimated from the measured data. This section details our technique for estimating the geometric model for the portions of the complete unknown axially symmetric surface associated with each of the measured fragments.

We start by fixing the fragment 3D data point constellation in arbitrary position in (x, y, z) 3D space, and then estimate the coordinate system $\{\mathbf{e}_r, \mathbf{e}_z\}$, or equivalently, the pot axis. A quick estimate is obtained using a generalization of the Hough transformation to 3D lines proposed in [8]. In [8], a linear method is detailed for computing the 3D line of minimum distance (in Plücker coordinates) to the linear complex, i.e., the set of 3D lines, defined by each measured surface point and its associated surface normal. Ideally this line is the pot central axis since, for an ideal surface of axial symmetry, the 3D line defined by each surface point and its associated local surface normal exactly intersects the axis of revolution. This estimate is further refined by the non-linear optimization specified in [7] which computes the axially-symmetric 3D algebraic surface of best fit for measured fragment data. This computation copies that in [6, 7]. The approach from this point on is different from those in [5, 4, 6, 7]. Hence we now have the pot axis with respect to the fixed fragment data. Then we Euclidean transform the fragment data such that it is in “standard position,” that is, the estimated fragment axis is aligned with the world coordinate z -axis, and the centroid of the fragment data lies on the world coordinate x -axis. In this coordinate system, we do have a potential estimated profile curve, namely that provided by the algebraic curve that specifies the estimated algebraic surface. However, we do not use this because it is a smoothed estimate of high frequency detail in the true profile curve. Hence, once we have the estimated fragment axis and the fragment data in “standard position,” we estimate the profile curve parameters \mathbf{P} using the fragment boundary curve. The fragment boundary curve is a simple closed 3D space-curve which, when projected in the (r, z) profile curve space is a complicated and often self-intersecting closed curve that traverses the unknown fragment profile curve twice; *once by passing clockwise* from one endpoint of the profile curve to the

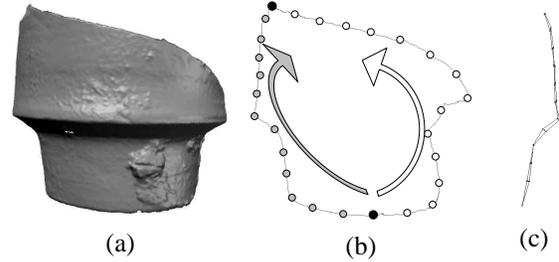


Figure 3: (a) A fragment of a Nabatean cup (dated to ~ 400 B.C.) cup excavated from Petra, Jordan (special thanks to Martha Joukowsky and the authors of [7] for access to this data). (b) Two possible paths traveling clockwise (in light gray) and counter-clockwise (in white) from one endpoint of the fragment profile to the other (in black). (c) The two profile curves generated by the traversals of the profile mentioned in (b) and their geometric mean : the estimated profile curve model.

other and *once by passing counter-clockwise* between the same two points (see fig 3). There are two important motivating factors to using the fragment boundary : (1) the boundary is guaranteed to traverse the two endpoints of the profile curve, and (2) for ideal axially-symmetric surfaces, the two profile curve parameterizations from the boundary will correspond to the maximum observed variation in the profile curve shape.

We then traverse the projected boundary parameterization of the profile curve clockwise and counter-clockwise, sampling the curve at constant intervals of arc-length for each traversal. This generates 2 sets of profile curve parameters which may (and often do) have different lengths, i.e., a different number of samples. The shortest of the two curves defines the total number of profile curve parameters and each point in \mathbf{P} is taken as the geometric average of a point on the shortest profile curve with the closest point found on longer profile curve. Note in the (r, θ, z) coordinate system the estimated pot model over the pot height interval pertinent to the measured fragment data is specified by $2N$ parameters \mathbf{P} , and the associated fragment data position which is a Euclidean transformation of the original fragment data position is specified by 4 parameters : the distance from the origin along the x -axis to the fragment data centroid, and the orientation of the fragment data constellation about its centroid, that is, its rotation angles about unit vectors parallel to the $x, y,$ and z axes.

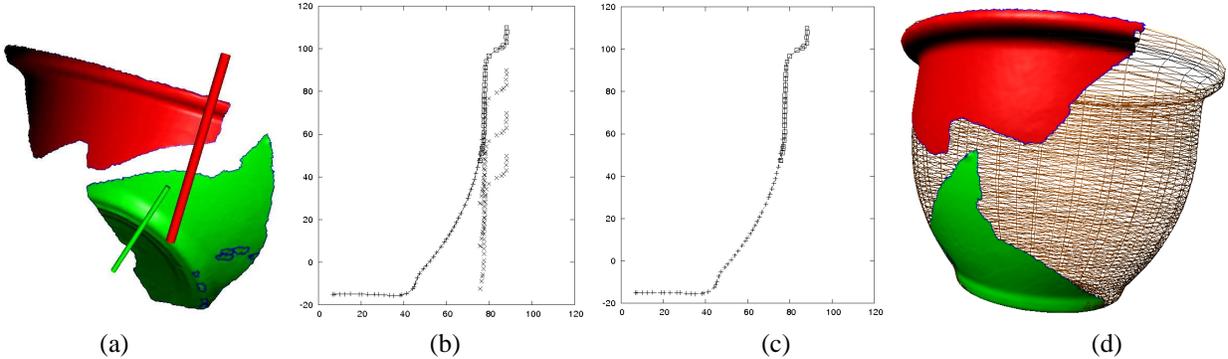


Figure 4: (a) Two fragments from a contemporary pot in arbitrary 3D orientation with their estimated mean axes rendered as cylinders (thanks to authors of [7] for access to this data). (b) We slide the profile-curve of one fragment with respect to the other to compute t_z , the best relative height alignment of the two fragments. (c) The aligned mean vectors for the fragment pair obtained by maximizing equation (1). The mean vectors of the resulting axis/profile-curves distributions are rendered as a set of connected 2D (r, z) points where the mean axis of symmetry is the vertical axis. Point sample locations are indicated with crosses '+' and boxes '□' for the two fragment models. Figure (d) shows the 3D axially symmetric surface generated from the mean vector as a wireframe mesh with the aligned fragment data superimposed. Of note here is that we may reconstruct the complete pot profile from fragments which may not share a common boundary as is the case in (d).

4 Obtaining Fragment Shape Distributions

In this section we augment our estimated fragment models by developing *shape distributions*; stochastic descriptions for each fragment that characterize the variability in the estimated fragment shape. We begin by applying the bootstrap method as detailed in [9] with a sample size of $B = 100$ bootstrap samples. Each bootstrap sample uses the same number of points and normals as the original fragment data set and is generated by a random re-sampling of the fragment point/normal data where a single data point may be selected more than once. For each of the bootstrap samples, an estimate of the geometric parameters, Θ , is computed. It is assumed that the resulting B estimates of the fragment axis and profile curve represent independent samples taken from a multivariate normal distribution with unknown mean and covariance. One may estimate the joint distribution of the axis and profile curve parameters by computing the estimated mean $\hat{\Theta} = \frac{1}{B} \sum_{b=1}^B \Theta_b$ and covariance $\hat{\Sigma} = \frac{1}{B-1} \sum_{b=1}^B (\Theta_b - \hat{\Theta})(\Theta_b - \hat{\Theta})^t$. These are sufficient to construct an approximation of the true distribution of the fragment surface parameters which are assumed to be normally distributed $\sim \mathcal{N}(\hat{\Theta}, \hat{\Sigma})$. In §5 we use this information to decide on the confidence the system can have in hypothesized fragment alignments.

In bootstrapping the profile curve parameters, we have chosen to use the variation of the projected fragment boundary as a source for characterizing random variations in the fragment profile curve. Another possible approach would proceed by estimating the variation of measured surface

points about an estimated curve in the (r, z) plane. However, this would discard the known correlation and ordering between measured points as offered by the connectivity of the fragment surface mesh and yield inaccurate results especially in high-curvature surface regions. Our approach is not completely consistent with the bootstrap procedure and will not capture all of the profile curve variations. Specifically, surface variations away from the fragment boundary will not be considered in this approach. An improvement would compute a set of random walks along the surface mesh from one profile-curve endpoint to the other. Each random walk would provide a random sample of a surface path spanning the profile curve parameters. These resulting profile curves could then be discretized to compute a more accurate distribution for the fragment profile curve parameters. We have chosen to use the boundary alone for computational speed but plan to investigate the possible benefits of using the more accurate method described.

5 Alignment Optimization

In this section we describe our approach for aligning groups of fragments using the stochastic model. We start by taking a pair of fragments from the set of K fragments and assume that the two fragments come from the same unknown pot and share some region of the unknown global profile curve. To align the fragments, we must determine what portion of their surfaces correspond and the best relative alignment between the pair given the computed correspondence.

To do so, we use our estimated parameters for each fragment as specified in §2 and §3 and the associated estimated

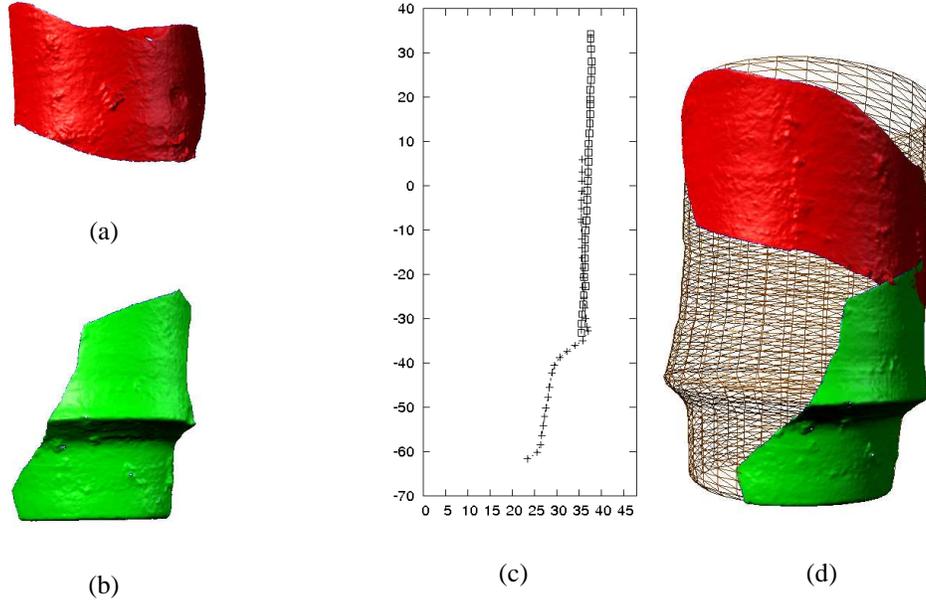


Figure 5: Figs (a-b) show surface measurements from archaeological fragments of the Nabatean cup from figure (3) which are to be aligned. Figure (c) shows the aligned mean vectors of the estimated axis/profile-curve distributions for the fragment pair. Figure (d) shows the 3D axially symmetric surface generated by merging the two distributions. It is shown as a wireframe mesh with the aligned fragment data superimposed.

Gaussian distributions for these parameters are obtained as specified in §3. Since the aforementioned fragment pair is assumed to have come from the same unknown pot and share some region of the unknown profile curve, the two estimated distributions share a common subspace of parameters. To combine the two distributions one must first solve for this subspace.

Before getting to the details, we point out that due to the cylindrical coordinate representation we have used, the alignment of two (or more) fragments becomes that of having two stochastic curves, each given in a common (r, z) coordinate system, and aligning these curves. The alignment is done by sliding one curve along its z -axis until the stochastic curves partially overlap such that the overlapped portions are stochastically similar, that is, such that the overlapped portions satisfy the relation of being two independent sample functions of the same 2D stochastic curve. Hence, we keep curve i fixed and translate curve j along the z -axis by an amount \mathbf{T}_{ji} . The position of the fragment measured data does not play a direct role here and can be ignored, because we do not make any use of how fragment data sets might touch and fit together.

Let $\{\Theta_i, \Theta_j\}$ denote the pair of parameter vectors for the fragments (i, j) which are to be aligned. The alignment algorithm proceeds by fixing one fragment, referred to as Θ_i in standard position as described in §3. The second fragment, referred to as Θ_j is translated along the z -axis within

the $(r, \theta = 0, z)$ plane subspace to maximize the joint probability of the two surfaces given the transformation value. Let \mathbf{T}_{ji} denote the translation that transforms fragment Θ_j with respect to Θ_i . Using our shape distributions from §4, the best surface alignment then corresponds to the Maximum Likelihood Estimate (MLE) of the translation parameter given the shape distributions for the fragment pair (1).

$$\hat{\mathbf{T}}_{ji} = \max_{\mathbf{T}_{ji}} p(\Theta_i, \Theta_j | \mathbf{T}_{ji}) \quad (1)$$

6 Model Regularization

The curve models adopted in §2 use an ordered sequence of (r, z) points to represent a 2D curve. In this section we discuss our resolution of two issues which result from this model choice :

1. Profile-curves for different fragments may have different parameterizations, i.e., we are matching discrete samples from a curve which may not lie at corresponding locations.
2. For smooth curve regions neighboring 2D curve samples may be very highly correlated. In this case, the covariance matrix estimate, i.e., $\hat{\Sigma}$ from §4, will be close to singular and numerically unstable for computation, i.e., inversion.

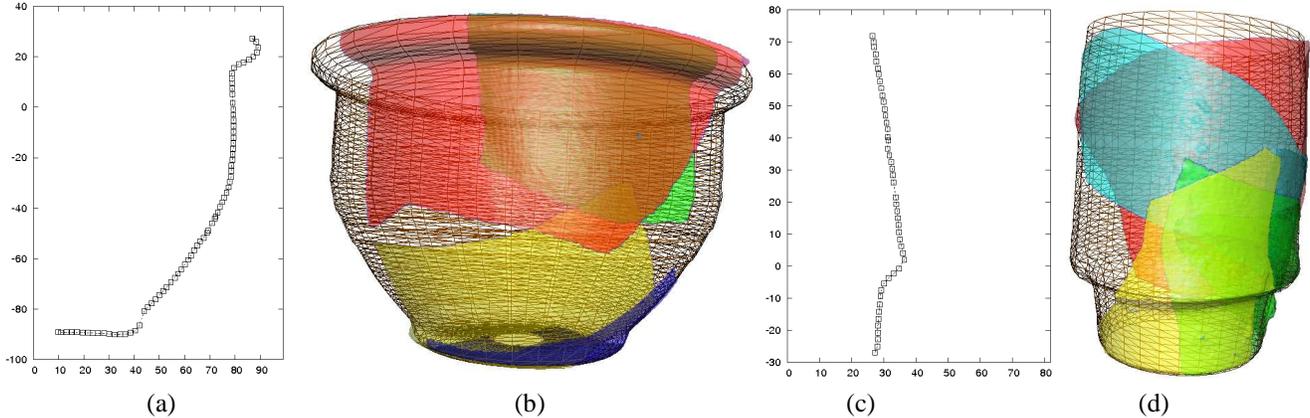


Figure 6: Results for more than 2 fragments : Figures (a) and (c) are mean vectors for the axis/profile-curve distributions generated by merging two additional surface models to the surfaces shown in figures (4) and (5). The 3D surfaces generated by revolving the mean profile-curve around the mean 3D axis of symmetry are shown as wireframe meshes with the aligned fragment data superimposed as semi-transparent regions with 4 different colors.

Issues (1) and (2) are surprisingly linked : One may address the parameterization problem by sampling the curves very densely, i.e., reducing the space between curve samples, Δs , such that the maximum parameterization error for a single matched sample is bounded above by $\frac{1}{2}\Delta s$ (where s is arc length along the fragment profile-curve). However, in smooth curve regions, i.e., regions well approximated by a straight line, as we sample more densely the correlation between neighboring curve samples will become extremely high. This will result in overfitting, i.e., allocating more parameters to the model than is necessary. The additional parameters will then represent the measurement noise and are essentially dependent random variables. Analytically, when two random variables in multi-dimensional covariance matrix are completely dependent the covariance matrix is singular. In practice, the covariance matrix $\hat{\Sigma}$ becomes close to singular and is not suitable for computation. One approach to this issue is dimensionality reduction, i.e., computation of second parameter set which spans a subspace of the chosen parameterization. The major problem here is that the *resulting subspace parameterizations are unique to each fragment* which introduces much analytical and computational complexity when matching parameterizations from two (or more) different fragments, i.e., dimensionality reduction makes matching *very* difficult for multiple fragment configurations.

We approach this problem by regularizing the covariance matrix using a non-linear variation of the well-known Tikhonov regularization [10] also referred to in statistical literature as ridge regression [11, 12]. This approach seeks to smooth the observed variances in the covariance matrix using a regularization function $\omega_\alpha(\lambda_i)$. This function improves the condition number of the matrix and makes the

computation of $\hat{\Sigma}^{-1}$ possible when parameters are highly correlated. Specifically, we compute the singular value decomposition of $\hat{\Sigma} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}'$ where $\mathbf{\Lambda}$ is a diagonal matrix of with positive eigenvalues λ_i . We then replace the eigenvalues according to the non-linear weighting (2) to generate a new diagonal matrix $\mathbf{\Lambda}_{new} = \log(10\mathbf{I} + \mathbf{\Lambda})$.

$$\omega_\alpha(\lambda_i) = \log(10 + \lambda_i) \quad (2)$$

We can then use the regularized matrix $\hat{\Sigma}_{biased} = \mathbf{U}\mathbf{\Lambda}_{new}\mathbf{V}'$ for stable numerical computations.

7 Merging Surface Estimates

In this section, we detail our technique for merging a matched pair of surface distributions. We start with two estimated surface distributions $p_1(\Theta_1)$ and $p_2(\Theta_2)$ which may be either individual fragment models as in §3 or a surface distribution obtained by merging a set of aligned fragment models. Let Θ denote the true parameter vector for the global profile curve. The estimated parameters Θ_1 and Θ_2 are profile curves consisting of (r, z) points which describe some portion of the complete profile. The estimation algorithm for computing the joint MLE $\hat{\mathbf{T}}_{21}$ from (1) provides a point-to-point correspondence, i.e., a shared region of the complete profile curve where we are matching (r, z) point parameters from each profile curve as described in §5. Let Θ_s denote these *shared parameters*, i.e., the common subspace of the two distributions which we must combine to merge our surface distributions. We add an index to the subscript $\Theta_{i,s}$ and $\Sigma_{i,s}$ to denote estimates of the shared mean vector and covariance matrix entries due to data from fragment i . Let $\Theta_{i,\neg s}$ denote the non-overlapping portions

of the mean vectors and $\Sigma_{i,\neg s}$ denote covariance matrix entries involving non-overlapping portions of the mean vectors. Given these new variables we can subdivide the covariance matrix for the two distributions. Σ_1 is divided into 4 regions as follows

$$\Sigma_1 = \begin{bmatrix} \Sigma_{1,\neg s} & \Sigma_{1,\neg s} \\ \Sigma_{1,\neg s} & \Sigma_{1,s} \end{bmatrix}$$

We re-order the rows and columns of Σ_2 according to the computed correspondence, such that shared parameters appear in the same order as adopted for Σ_1 , i.e.,

$$\Sigma_2 = \begin{bmatrix} \Sigma_{2,s} & \Sigma_{2,\neg s} \\ \Sigma_{2,\neg s} & \Sigma_{2,\neg s} \end{bmatrix}$$

We may then compute the new mean vector and covariance matrix of $p_{1,2}(\Theta_{1,2} | \mathbf{T} = \hat{\mathbf{T}}_{21})$ by concatenating the vectors $\Theta_{1,2} = \begin{bmatrix} \Theta_{1,\neg s} & \hat{\Theta}_s & \Theta_{2,\neg s} \end{bmatrix}'$ where equation (3) is applied to compute $\hat{\Theta}_s$ as specified in basic pattern recognition texts such as [13]. Likewise the covariance matrix is specified as :

$$\Sigma_{1,2} = \begin{bmatrix} \Sigma_{1,\neg s} & \Sigma_{1,\neg s} & 0 \\ \Sigma_{1,\neg s} & \hat{\Sigma}_{ss} & \Sigma_{2,\neg s} \\ 0 & \Sigma_{2,\neg s} & \Sigma_{2,\neg s} \end{bmatrix}$$

$$\begin{aligned} \hat{\Sigma}_{ss} &= \Sigma_{1,s}(\Sigma_{1,s} + \Sigma_{2,s})^{-1}\Sigma_{2,s} \\ \hat{\Theta}_s &= \Sigma_{2,s}(\Sigma_{1,s} + \Sigma_{2,s})^{-1}\Theta_1 + \Sigma_{1,s}(\Sigma_{1,s} + \Sigma_{2,s})^{-1}\Theta_2 \end{aligned} \quad (3)$$

8 Results

We present results which reconstruct the profile of 2 pots : (1) a contemporary mass-produced pot broken for experimentation (see fig. 4a) and (2) an Nabatean cup from approximately 400 B.C. uncovered in Petra, Jordan (see fig. 5a-b). In each case, we provide the system a set of pieces which come from a single pot. Each profile curve was sampled at constant intervals of arc length where $\Delta s = 2mm$. we used the constant of 6 mm. as the arc-length interval between sampled profile curve points as described in §2. In figures (4) and (5) the system optimizes (1) to merge a single fragment pair. In figure (6) the system The algorithm described above was implemented in Java on a 1.6GHz computer and the solutions to each alignment problem took approximately 1 second to compute the solution to each merge, i.e., for figures (4) and (5) run time was approximately 1 second and for figure (6) run time was approximately 3 seconds for both solutions with considerable opportunity for optimization still available. This is an improvement of an order of magnitude over past methods [6].

9 Conclusion

This paper introduces a new approach for 2D curve/3D axially-symmetric surface modeling and constrained matching for the purpose estimating the unknown shape of the global pot surface. Measured fragment data is used to generate shape distributions for each fragment which characterizes both the shape and shape variability of each fragment. Fragments are then matched together by computing the MLE of the transformation parameters given the joint shape distribution of the matched fragments. The approach is applicable to an arbitrarily large set of fragments and alignment results are illustrated for pairs and quadruplets of fragments. The approach provides a compact shape representation for each fragment consisting of 20-50 parameters for a fragment. The compact representation greatly simplifies the difficult task of matching 3D surfaces with 10k-200k measurements and requires much less computation without a significant loss in accuracy.

References

- [1] H. Wolfson, E. Schonberg, A. Kalvin, and Y. Lapidan, "Solving Jigsaw Puzzles Using Computer Vision," *Annals of Operations Research*, vol. 12, pp. 51–64, 1988.
- [2] J. Stolfi and H.C. Leitão, "A multiscale method for the reassembly of two-dimensional fragmented objects," *PAMI*, 2002.
- [3] J. McBride, "Archaeological Fragment Re-Assembly Using Curve Matching," M.S. thesis, Brown University, September 2002.
- [4] G. Papaioannou, E-A. Karabassi, and T. Theoharis, "Reconstruction of three-dimensional objects through matching of their parts," *PAMI*, 2002.
- [5] W. Kong and B. B. Kimia, "On solving 2D and 3D puzzles using curve matching," in *Proc. of CVPR*, Hawaii, USA, December 2001, IEEE, Computer Society.
- [6] A. Willis, *Stochastic 3D Geometric Models for Classification, Deformation, and Estimation*, Ph.D. thesis, Brown University, May 2004.
- [7] A. Willis, D. Cooper, et al., "Bayesian Pot-Assembly from Fragments as Problems in Perceptual-Grouping and Geometric-Learning," in *ICPR*, 2002, vol. III, pp. 297–302.
- [8] H. Pottmann, M. Peternell, and B. Ravani, "An introduction to line geometry with applications," *Computer-Aided Design*, vol. 31, pp. 3–16, 1999.

- [9] B. Efron and R. Tibshirani, *An Introduction to the Bootstrap*, Chapman-Hall, 1993.
- [10] A. N. Tikhonov and Arsenin V. Y., *Solutions to Ill-Posed Problems*, Winston & Sons, 1977.
- [11] D. W. Marquardt, "Generalized inverses, ridge regression, biased linear estimate and nonlinear estimation," *Technometrics*, vol. 12, pp. 591–612, 1970.
- [12] A.E. Hoerl and R.W. Kennard, "Ridge regression : Biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, pp. 55–57, 1970.
- [13] R. Duda, P. Hart, and D. Stork, *Pattern Classification*, John Wiley & Sons, 2nd edition, 2001.