

Axially Symmetric 3D Pots Configuration System using Axis of Symmetry and Break Curve

Kilho Son
Brown University
Providence, RI

kilho_son@brown.edu

Eduardo B. Almeida
Brown University
Providence, RI

eduardo_almeida@brown.edu

David B. Cooper
Brown University
Providence, RI

david_cooper@brown.edu

Abstract

This paper introduces a novel approach for reassembling pot sherds found at archaeological excavation sites, for the purpose of reconstructing clay pots that had been made on a wheel. These pots and the sherds into which they have broken are axially symmetric. The reassembly process can be viewed as 3D puzzle solving or generalized cylinder learning from broken fragments. The estimation exploits both local and semi-global geometric structure, thus making it a fundamental problem of geometry estimation from noisy fragments in computer vision and pattern recognition. The data used are densely digitized 3D laser scans of each fragment's outer surface. The proposed reassembly system is automatic and functions when the pile of available fragments is from one or multiple pots, and even when pieces are missing from any pot. The geometric structure used are curves on the pot along which the surface had broken and the silhouette of a pot with respect to an axis, called axis-profile curve (APC). For reassembling multiple pots with or without missing pieces, our algorithm estimates the APC from each fragment, then reassembles into configurations the ones having distinctive APC. Further growth of configurations is based on adding remaining fragments such that their APC and break curves are consistent with those of a configuration. The method is novel, more robust and handles the largest numbers of fragments to date.

1. Introduction

History can be uncovered by discovering ancient artifacts, such as pots and tools. Meaningful historical information about human society is inaccessible when artifacts are found shattered into many pieces. The shape of a single fragment, also known as *sherd*, contains negligible information when compared to the complete shape. If the pieces are reassembled, the hidden historical records can be retrieved. The large majority of ceramic objects at archaeological sites (pots, bowls, plates, water jugs etc.) are axially symmetric,

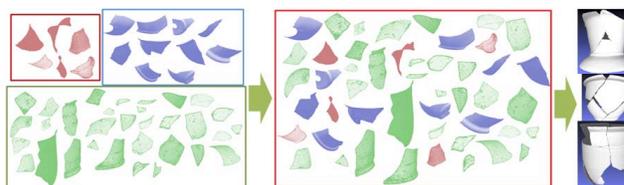


Figure 1: The system configures 3 pots from mixed and unorganized fragments.

and these objects are key to understanding how the community at the time lived. However, the fragments are commonly found mixed in very large numbers as a giant 3D puzzle. Manually reconstructing artifacts from hundreds of mixed broken pieces is a very tedious time consuming process. The demand for an automatic reconstruction framework to help archaeologists is evident and motivates the research in this topic.

The problem described in this paper can be viewed as a 3D puzzle problem [16, 15, 6, 7] dealing with a most important incompletely solved problem in archeology. Puzzle problems have been developing for many years starting from Freeman and Garder in 1964 [2] solving jigsaw puzzles using 2D shape information. While these early efforts focused more on matching shape of smaller 2D puzzles [2, 4], the recent advances tackle very large square-piece jigsaw puzzles, where shape information is missing, but piece alignment space is discrete and matching is based exclusively on image content [1, 12, 3]. The nature of reassembling 3D fragments (e.g. [14, 11, 6]) is in general more complex, since the alignment search is in a continuous space of 3D rotations and translations.

Regarding general assembly procedures, solutions typically consist of two steps, local compatibility of pieces and global strategy for puzzle assembly. Various local measures are used, such as the ones based on curve matching [10, 17], junctions [10], features [4] and texture information. Greedy algorithms [4, 10, 12], iterative relaxation methods [17] and graphical models based on belief propagation [1] have been

studied for global configuration of a puzzle. In [11], 3D complementary fractured surfaces are matched using simulated annealing applied to depth images. Huang *et al.* [6] detect break surfaces and reestablish their connections by corresponding robust geometric descriptors and avoiding fragment intersection.

Pot assembly is more challenging if no texture information is given, as the problem this paper tackles. A solution is presented with focus on ceramic fragments from axially symmetric pottery (solids of revolution). The proposed method is designed to completely configure multiple pots from arbitrary fragments based on the knowledge of their outer surface. The system pipeline is unique in its use of axial information, which has so far not received much attention for reassembly purposes, despite its inherently strong cues. This paper method significantly improves the state-of-the-art [14], which is driven primarily by matching break curves, secondarily by using APC and is only semi-automatic. Furthermore, it is shown that the system is effective on real fragments from archaeological excavation since it is robust to noise and perturbations such as break curve chipping, erosion and to 3D noise in input surfaces.

The main technical contributions of this paper are:

- Reassembly of larger numbers of pieces per pot than previously possible,
- Fully automated and simultaneous reassembly of more than one pot,
- Use of the global constraint that all fragments share the same axis significantly reduces computations and increases robustness to noise,
- The noise present in scanned pottery data used in this work was previously unmanageable.

1.1. Problem statement

Ceramic pots of interest have been made on a wheel and are thus axially symmetric. These are thin wall pots and their geometry is specified by their silhouettes with respect to an axis. Hence, the intersection of a plane perpendicular to the axis with the outer pot surface is a circle. The silhouette as a function of height along the axis can be single-valued or multi-valued. Denote the axis-silhouette combination as the axis/profile curve (APC). The pot breaks into sherds (fragments) along curves on the outer surface. Denote these curves as *break curves*. Hence, for each fragment we refer to the outer-surface curve that specifies the fragment shape as its break curve. Typically, a pot at an archaeological excavation site is broken into anywhere from a few fragments to 100-200 fragments. 20-30 is a representative number. The goal is to work with a pile of sherds and reconstruct the one or more pots that gave rise to these sherds. The sherds present for any one pot may not be complete. We do not use surface albedo patterns in the reconstruction, nor do we use pot wall thickness. We use only x,y,z measure-

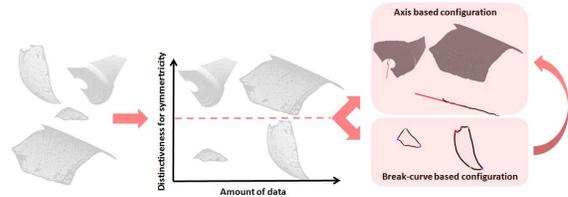


Figure 2: System overview

ments of the fragment outer surface and the fragment break curve, where the break curve estimates are computed from the outer surface measurements. Typically, 2,000 - 10,000 3D points are measured per fragment. These measurements are noisy, especially those of the break curve. The break curve measurements are noisy for two reasons. First, two fragments may not have broken in a way where the break surface is perpendicular to the pot outer surface. Thus, the break curve location is not well defined. Second, if great care is not exercised in estimating a break curve from the laser-measured surface points, the break curve estimate will be highly variable, and this is the case with the data used in this paper. *The reason we do not use surface albedo or surface thickness information is because our interest is seeing what can be done by measurements of just surface and break curve geometries. The other information of course could be very useful, and can be used in a more complete treatment of pot sherds reassembly.* Useable reassembly of fragments for a pot does not require reassembling all of the original pieces from the broken pot. Of major use is estimating the APC for the original pot. This conveys all of the geometry for the original pot, and that is the goal reported on in this paper.

1.2. System overview

First, the system estimates the axis of symmetry for each fragment using circle templates [9], i.e., cutting the fragments with planes normal to this axis the intersections are circles (Section 2). The fragments are classified as to whether the fragment data contains enough and distinctive information for estimating the axis of symmetry (Section 3). If fragments' estimated axes are reliable, fragment reassembly is based on their symmetry property and break curve, otherwise is based solely on break curves. This procedure is essential, because if the axis estimation is unreliable the system fails when using it.

Second, for fragments with reliable axis estimation, the system refines their axis by matching fragments in pairs and re-estimating the common axis for each pair. Given pair matching candidates, the pots configuration cost is formulated as a quadratic form, which penalizes deviations from global symmetry, constrained with geometrical consistencies, and can be approximately minimized by a spectral technique [8] (Section 4).

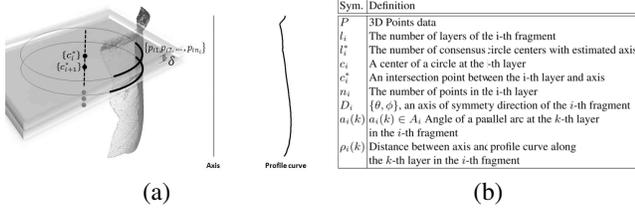


Figure 3: (a) Axis of symmetry estimation method and axis/profile curve (APC); (b) Notation.

Finally, for the fragments whose axis estimation is unreliable, the system uses a crude feature matching method based on only break curve features. The system represents a small patch of break curve with a weighted histogram of curvature orientations. The break curve matching results are verified with previously obtained axis information (Section 5).

2. Axis of symmetry estimation

Configuring the relative position between two broken objects in 3D is a challenging problem. Many papers present methods to solve this problem by registering break surfaces [6, 11]. In order to optimally align 3D break curves from noisy and possibly incomplete 3D measurement data of ancient artifacts outer surface, strong cues are taken into consideration. This paper tackles the problem with a symmetry property, by first estimating an *axis of symmetry*. Axis estimation needs to be very accurate since it is used throughout the entire framework. This section describes how to estimate the relative position and direction of the axis.

Previous solutions can be roughly classified into two categories, either exploiting surface normal vectors which point toward the axis [13] or making use of circle templates [9, 5]. Willis [13] uses bootstrap algorithm to enhance robustness to noise caused by geometric distortions and measurement errors. Based on experiments, Mara and Sablatnig [9] proposed a circle template method claiming that the normal vectors are sensitive to noise and fragment shapes. It is assumed that surface points intersecting a thick plane perpendicular to the true axis should be in circular shape. Both algorithms have different benefits. Willis' algorithm provide accurate axis estimation, in spite of perturbations, when the broken pieces are large enough, otherwise it fails, while Mara's approach can also handle relatively small pieces at lower accuracy. The proposed archaeological application requires assembly of even smaller pieces. Hence, a method based on circle template is proposed for estimating axis of symmetry.

A given fragment \mathcal{F} is represented as a collection of points $P = \{p_{ij} \in \mathcal{F}\}$ where p_{ij} is j -th point in the i -th layer. Given an arbitrary 3D direction, the 3D points are divided into multiple planar layers of thickness $\delta = 0.5$, which is 1.1 times the scanner resolution, and several cir-

cles are fit to the data using a least squares method. Then a 3D line is fit to the estimated circles centers with RANSAC. The 3D line must be the fragment axis of symmetry by definition of SOR. This constrains the axis location for the given orientation. Among every possible hypothetical axis direction $D = \{\phi, \theta\}$ with steps of 5° , consisting of two angles, the optimal axis is estimated by inexpensive brute-force search based on the cost function

$$E_a(D|P) = \frac{\sum_{i=1}^l n_i \sum_{j=1}^{n_i} \left(\|p_{ij} - c_i^*\| - \mu_i \right)^\beta}{\sum_{i=1}^l n_i} - \alpha \log l^*, \quad (1)$$

where c_i^* is crossing point between the i -th layer and the estimated axis, n_i is number of 3D points within the i -th layer, $\mu_i = \frac{1}{n_i} \sum_{j=1}^{n_i} \|p_{ij} - c_i^*\|$ is the i -th layer average estimated radius, l^* is number of consensus circle centers with the estimated 3D axis, and l is total number of layers generated by an axis hypothesis (see Fig. 3).

The first term in equation (1) represents the symmetry cost of axis D . The parameter β should be smaller, if the noise on p_{ij} is significant. The second term applies minimum description length (MDL) principle. One of the important properties of the MDL method is that a safeguard is provided against overfitting by a trade-off between the complexities of the hypothesis and of the data given the hypothesis. The rationale is that the estimated axis can be supported by a small number of circle centers, therefore the estimation may be overfitted. α is a weight between the cost function terms. The parameters are empirically determined as $\alpha = 1$, $\beta = 2$. The optimal direction estimation, D^* of the axis of symmetry is

$$D^* = \{\phi^*, \theta^*\} = \arg \min_{\theta, \phi} E_a(\phi, \theta|P). \quad (2)$$

The optimal axis position is estimated by fitting a line to the circle centers associated with the optimal direction D^* .

3. Fragment types

This paper propose a reasonable way for estimating axis of symmetry by minimizing the cost function $E_a(D|P)$. However, this method does not always return the true axis. For example, for a flat or sphere-shaped fragment, the minimum of $E_a(D|P)$ is achieved for infinitely many directions. If the axis estimation is inaccurate or ambiguous, the pipeline reliability may be compromised. To avoid this risk, the system classifies the fragments as four types and adopts an approach based on the type. The fragment classification criteria are: *sufficiency* and *distinctiveness* for axis estimation. Sufficiency can be measured by the cost of $E_a(D^*|P)$ where D^* is given in equation (2). Smaller values correspond to axis reliably estimated with support from many thick layers. Analogously, higher values means little support and less fidelity.

To measure distinctiveness of the data, the variation $O(\phi^*, \theta^* | P)$ of E_a in the neighborhood of global minimum D^* is analyzed. If the axis estimation is reliable, the cost value increases fast as the axis direction is slightly changed and vice versa. The variation is given by

$$O(\phi^*, \theta^* | P) = (E_a(\phi^*, \theta^* | P) - E_a(\phi^* + \Delta\phi, \theta^* + \Delta\theta | P))^2. \quad (3)$$

This is approximated by a first order Taylor expansion

$$E_a(\phi^* + \Delta\phi, \theta^* + \Delta\theta) = E_a(\phi^*, \theta^*) + \left[\frac{\partial E_a(\phi^*, \theta^*)}{\partial \phi} \frac{\partial E_a(\phi^*, \theta^*)}{\partial \theta} \right] \begin{bmatrix} \Delta\phi \\ \Delta\theta \end{bmatrix}. \quad (4)$$

P is omitted for notation simplicity. By substituting equation (4) on equation (3),

$$O(\phi^*, \theta^* | P) = \left(\left[\frac{\partial E_a(\phi^*, \theta^*)}{\partial \phi} \frac{\partial E_a(\phi^*, \theta^*)}{\partial \theta} \right] \begin{bmatrix} \Delta\phi \\ \Delta\theta \end{bmatrix} \right)^2 = [\Delta\phi \Delta\theta] \begin{bmatrix} \frac{(E_a(\phi^*, \theta^*))^2}{E_a(\phi^*, \theta^*) \frac{\partial E_a(\phi^*, \theta^*)}{\partial \phi}} & \frac{E_a(\phi^*, \theta^*) \frac{\partial E_a(\phi^*, \theta^*)}{\partial \theta}}{(E_a(\phi^*, \theta^*))^2} \end{bmatrix} \begin{bmatrix} \Delta\phi \\ \Delta\theta \end{bmatrix} = [\Delta\phi \Delta\theta] \Lambda(\phi^*, \theta^*) \begin{bmatrix} \Delta\phi \\ \Delta\theta \end{bmatrix}. \quad (5)$$

The metric captures the variability of the cost function according to the variables ϕ and θ . Moreover, the amount of variability of each direction can be analyzed by calculating the eigenvalues λ_1, λ_2 of $\Lambda(\phi^*, \theta^*)$ and their corresponding eigenvectors φ_1, φ_2 . A measure, inspired by corner detection, is defined for the axis estimation confidence as

$$O_f(P) = \lambda_1 \lambda_2 + \gamma(\lambda_1 + \lambda_2)^2, \quad (6)$$

where $\Delta\theta = \Delta\phi = 10$ degrees and γ is set to 0.4 empirically.

4. Matching based on axis estimation

In this section, a reconstruction approach for fragments carrying a distinctive symmetry property is described. The problem bears many difficulties due to noise on 3D measurements and large search space. In 3D, each fragment has 6 degrees of freedom (DOF) for rigid body motion.

The system makes use of the symmetry property to mitigate these difficulties in two steps: local matching and global configuration, as many previous puzzle solvers. First, the system defines a cost function for matching two fragments which represent how well the axis/profile curve (APC) and the break curve matched. Relative position between two fragments is estimated by minimizing the cost function. The search space is reduced by aligning two axes of symmetry, D_j^* and D_ℓ^* , which are accurately re-estimated as matching the two pieces simultaneously. This procedure reduces search space from 6 DOF (for a pair of fragments)

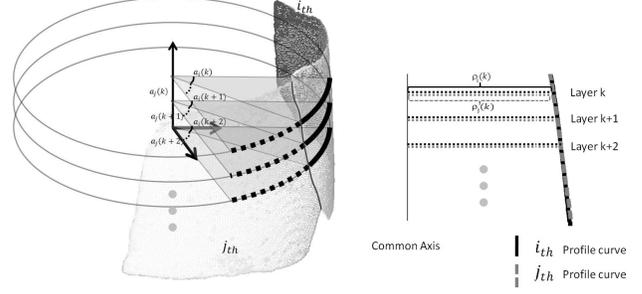


Figure 4: *Left*: two perfectly aligned sherds and an illustration of the constant angle sum constraint: the intersection of the sherds with planes perpendicular to their axis of symmetry defines circle arcs layers on each sherd, with angles $a_j(k)$ for the j -th sherd at the k -th layer. Clearly $a_j(k) + a_i(k)$ is constant. *Right*: sherds respective APCs.

to 1 DOF (relative position of 2 pieces along the axis of symmetry). Second, the system formulates the problem as minimizing a constrained quadratic cost function for semi-global configuration. The cost function captures symmetry constraints between candidate pairs and is minimized by a spectral method [8]. The symmetry constraint not only strongly improves pairwise matching cost, but also models global consistency between fragments not necessarily adjacent to each other, as far as they share an APC. The system can afford to use this extremely useful axis of symmetry property because its estimation is accurate enough.

4.1. Local Pair Matching

For matching 2 fragments, the proposed system basically uses an energy minimization scheme. The energy cost includes symmetry property and break curve information of the fragments. If two fragments are correctly aligned, the pair should satisfy two constraints: constant angle sum and APC alignment:

Constant angle sum (C1): this conception is introduced by [4] for solving 2D jigsaw puzzles by shape. When pieces fit perfectly, the distance between certain points must be constant. This constraint is simple but powerful. Figure 4 proposes a 3D extension over the pot surface.

Profile curve alignment (C2): a solid of revolution can be virtually reconstructed from its APC, which naturally constrain the position of pot sherds (see Figure 4). This additional constraint greatly improves matching performance by reducing search from 6 DOF to 1 (relative translation of two fragments along an axis).

Let $A_i = \{a_i(1), a_i(2), \dots, a_i(k), \dots, a_i(l_i)\}$ be arc angle ranges of the i -th fragment in k -th layer as described in Figure 4. These arcs are denoted *parallel arcs* (they never intersect). $\rho_i(k)$ is the distance between axis of symmetry and profile curve along the k -th layer in i -th fragments, as illustrated in Figure 4. $\rho_i(k)$ is denoted *layer radius*. The matching cost is formulated as a function between the i -th

and j -th pieces with the above two constraints:

$$E_m(\zeta_{ij}, D_i, D_j | P_i, P_j) = \mathbf{1}(n_o \geq \kappa) E_c(\zeta_{ij}, D_i, D_j | P_i, P_j) + w_1(E_p(\zeta_{ij}, D_i, D_j | P_i, P_j) + E_a(D_i | P_i) + E_a(D_j | P_j)), \quad (7)$$

where $\mathbf{1}()$ is an indicator function and

$$n_o = \sum_{k \in \Omega} \mathbf{1}(|a_i(k) + a_j(k - \zeta_{ij}) - \tau| < \eta),$$

$$\tau = \max_s (a_i(s) + a_j(s - \zeta_{i,j})),$$

where ζ_{ij} is relative translation between two i, j pieces along the common axis. The matching cost function includes two terms, E_c and E_p , which represent constraints **C1** and **C2** respectively. Ω has all common layers of two fragments. n_o represents the number of layers which agree with constraints **C1** with tolerance η .

$$E_c(\zeta_{ij}, D_i, D_j | P_i, P_j) = -e^{w_2(\eta - \mu_g)} - e^{w_3(n_o - \kappa)} + 1, \quad (8)$$

where,

$$\mu_g = \frac{\sum_{k \in \Omega} \mathbf{1}(\Upsilon_{ij}^k < \eta) \Upsilon_{ij}^k}{\sum_{k \in \Omega} \mathbf{1}(\Upsilon_{ij}^k < \eta)},$$

$$\Upsilon_{ij}^k = \tau - (a_i(k) + a_j(k - \zeta_{ij})),$$

$$E_p(\zeta_{i,j}, D_i, D_j | P_i, P_j) = \frac{1}{T_o} \sum_{k \in \Omega} |\rho_i(k) - \rho_j(k - \zeta)| + \gamma \left| \frac{\Delta \rho_i(k)}{\Delta k} - \frac{\Delta \rho_j(k - \zeta)}{\Delta k} \right|, \quad (9)$$

where T_o is total number of common layers of two fragments. $w_1 = 1$, $w_2 = 1$, $w_3 = \frac{2}{T_o}$, $\eta = 2/3$, $\gamma = 1/2$ and κ is $0.6T_o$. The E_m cost function has 5 DOF, 4 for D_i, D_j and 1 for relative position between fragments along the shared axis. E_m search space is reduced to 1 DOF by initializing the axis of symmetry as described in Section 2. The system re-estimates the axis of symmetry of both fragments by minimizing E_m to match two fragments robustly with a particle filter. The Gaussian distribution is applied to generate 10 axes samples for each fragment 3 times iteratively. Initial means and variances are respectively estimated by $\arg \min_{\phi, \theta} E_a(\phi, \theta)$ and from $\Lambda(\phi, \theta)$ in previous steps. The variance are the same for all iterations. $\zeta_{i,j}$ is learned by a brute-force method. A match of i -th and j -th pieces are considered as a candidate of a true match if $\mathbf{1}(n_o(\zeta_{ij}^*, D_i^*, D_j^*) \geq \kappa)$ is 1 where $\{\zeta_{ij}^*, D_i^*, D_j^*\} = \arg \min_{\zeta_{ij}, D_i, D_j} E_m(\zeta_{ij}, D_i, D_j | P_i, P_j)$. It means that once a pair match satisfies **C1** loosely, the algorithm considers it as a candidate pair match.

4.2. Semi-Global Configuration

This section will describe a semi-global configuration given candidate matches. Although the previous step uses

two constraints for matching, the results have some false alarms since the matching criterion is not strict. One of the strong points of axis of symmetry property is that although the fragments are not adjacent, they interact with each other. That is, if they are from single artifact and overlap at part of the profile curve, the their profile curve is the same at the overlapped region. This constraint is used to make a cost function for semi-global configuration. Let $X = \{x_1, x_2, \dots, x_L\} \in \{0, 1\}^L$ be a binary vector, where L is total number of pair candidates and $x_k = 1$ indicates that the k -th pair candidate is a true pair, otherwise it is not. The final goal is to find the minimizer X such that

$$X = \arg \min_{X \in \{0,1\}^L} \sum_{i=1}^L x_i E_l(c_i) + \sum_{i=1}^L \sum_{j=1, j \neq i}^L x_i x_j E_g(c_i, c_j),$$

where $C = \{c_1, c_2, \dots, c_L\}$ is a set of candidate pairs and $c_k = \{P_i, P_j\}$. The semi-global cost function is

$$X = \arg \min_{X \in \{0,1\}^L} X^t \mathbf{M} X, \quad (10)$$

where,

$$\mathbf{M} = \begin{bmatrix} E_l(c_1) & E_g(c_1, c_2) & \dots & E_g(c_1, c_L) \\ E_g(c_2, c_1) & E_l(c_2) & \dots & E_g(c_2, c_L) \\ \vdots & \vdots & \ddots & \vdots \\ E_g(c_L, c_1) & E_g(c_L, c_2) & \dots & E_l(c_L) \end{bmatrix},$$

$E_l(c_k) = E_l(P_i P_j) = \min_{\zeta_{i,j}, D_i, D_j} E_m(\zeta_{i,j}, D_i, D_j | P_i, P_j)$ and $E_g(c_i, c_j) = \min_{\zeta_{i,j}} E_p(\zeta_{i,j} | c_i, c_j)$.

The vector X is constrained so that relative positions of all fragments satisfy geometrical consistence such as a non-overlap between fragments and non-duplication of the shards. True pairs are found by estimating X which minimize the semi-global cost function with a spectral technique [8]. The configurations are determined by the estimated X .

5. Break Curve Matching

Certain fragments lack distinctive symmetry property since they are relatively small, planar or spherical. In this section, a method is introduced to deal with these sherds that cannot be treated with previous approach. A reasonable alternative cue for matching fragments is found at salient features on their break curves. However, after chipping and erosion, local features such as curvature and torsion between corresponding points on matching curves can be very noisy, so direct comparison is prohibitive.

To solve this problem, a robust matching approach based mainly on break curves is designed to involve feature points. First, the break curves are smoothed and uniformly sampled. Second, multiple scale *Horoscope* (Histogram Of Relative OSculating Circle and Outer surface tangent Planes orientations Estimates) are computed at each break curve

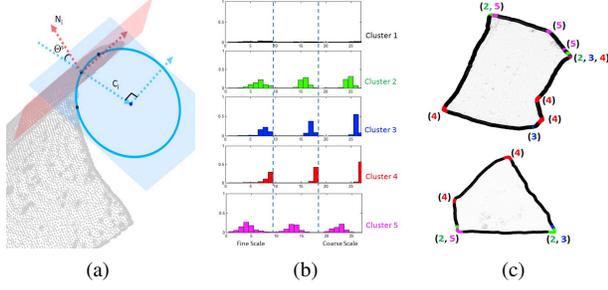


Figure 5: (a) Osculating circle and surface normal vector; (b) weighted orientation histogram for each cluster; (c) break curves and their cluster labels.

point. The histograms are used to detect salient corners and cluster them based on similarity. Finally, distinctive corner clusters are matched and the APC is used for sherd alignment refinement and to discard false matches. This procedure drastically reduces curve matching complexity from a huge combinatorial search among all possible pairs of curve segments to a few candidates. Detected corners identify reasonable break curve endpoints which are clustered and help to pinpoint the candidates segments for matching.

5.1. Break Curve Smoothing

Next section describes how to construct a multiscale *Horoscope* along break curve points. Prior to computing *Horoscope* features, some pre-processing is required to smooth all break curves and sample them at fixed spatial resolution. Thus, the three coordinates of 3D curve are independently convolved with a $[1\ 2\ 3\ 5\ 3\ 2\ 1]/17$ kernel and then linearly interpolated at uniform distances. The uniform distance interpolation is performed by starting from a given curve point P_i , intersect the curve with a small sphere S_g of fixed radius R_g centered at P_i to obtain the next point P_{i+1} , then repeat. R_g is chosen such that roughly 3 nearest neighbors of P_i are enclosed by S_g . Therefore, distance in this context is not arc length, but Euclidean distance, a more appropriate choice to deal with bumps on the break curve. This sequence of points in a curve is denoted as $P_1^n = \{P_1, \dots, P_n\}$ and their spatial resolution is constant regardless of how the sherds were scanned.

5.2. Horoscope features

A descriptor Θ at points $P_i \in P_1^n$, invariant to proper rigid transformation, is given in equations (11) and (12). $\Theta(P_i)$ depends on the orientation between the osculating circle plane and the tangent plane at P_i .

$$C^s(P_i) = \frac{(\hat{P}_{i+1}^s - P_i) - (P_i - \hat{P}_{i-1}^s)}{R_s^2} \quad (11)$$

$$\Theta^s(P_i) = \cos^{-1} \left| \frac{C_i \cdot N_i}{\|C_i\| \|N_i\|} \right| \quad (12)$$

The index s refers to scale. $C^s(P_i)$ is the estimated 3D curve normal vector at P_i , pointing at the center of its osculating circle. \hat{P}_{i+1}^s and \hat{P}_{i-1}^s are intersections between the break curve and a sphere of radius $R_s > R_g$ centered at P_i . R_s is the scale parameter. The norm of $C^s(P_i)$ represents curvature. N_i is the normal of the sherd tangent plane at P_i . $\Theta^s(P_i) \in [0, \frac{\pi}{2}]$ is the *Horoscope* angle between $C^s(P_i)$ and N_i . Dependencies in P_i on equation (12) are omitted for simplicity of notation. The use of $\hat{P}_{i\pm 1}^s$ instead of $P_{i\pm 1}$ in eq. (11) is for robust curvature estimation with respect to noise still present on the sequence P_1^n .

A weighted orientation histogram $\Gamma^s(P_i)$ of Θ^s is defined as follows: 11 data points of a small arc $\{\Theta^s(P_{i-5}), \dots, \Theta^s(P_{i+5})\}$ are distributed on the interval $[0, \frac{\pi}{2}]$, quantized into 9 bins. For each data point ξ , its count is weighted by $\|C^s(\xi)\|$. The histogram $\Gamma^s(P_i)$ has 9 curvature-weighted counts. After selecting 3 scales, $\{s_1, s_2, s_3\}$, the counts are concatenated into a 27-d vector $\hat{\Gamma}(P_i)$, the multiscale *Horoscope* feature at a point P_i .

This rather unusual feature is inspired by insights in scene classification, object recognition, and image and text retrieval, where visual words are introduced. The main idea of visual vocabulary is that they quantize features by clustering local robust descriptors. The *Horoscope* is a robust multiscale measure of local oriented curvatures which loosely describe noisy corner points.

5.3. Clustering and Matching

This section describes sherd matching and alignment based on break curve feature. For all available sherds, the *Horoscope* feature $\hat{\Gamma}$ is computed in all of their points. The features from all sherds are then clustered using k -means ($k = 5$). The outcome has essentially 4 distinct clusters corresponding to high curvature regions. An enormous cluster of nearly straight lines (Figure 5c: black curves) is disregarded. Matching is based solely on the high curvature regions clusters (see Figure 5c color-coded corners). In general, break curves corners are piece-wise constant sequences of the distinctive 4 clusters. A connected group of such sequences, in between the disregarded low curvature points, is denoted a *horoscope corner*, or simply *h-corner*.

In order to reduce matching complexity, candidate fragments are selected based on the *h-corners*. It is assumed that *h-corners* on two matching sherds share at least one common type of cluster. For a pair of candidates, the sherds can be connected at the *h-corners* endpoints. This process is very robust to noisy data. Iterative closest point (ICP) algorithm is used to compute a rigid transformation required for alignment. If candidates are genuine matches, the (noisy) estimated surface normals will also align and the sherds are considered true matches. The surface normal alignment is only used to validate the ICP result. To obtain a better initialization for ICP, each arc is temporarily approximated by

a straight line connecting its endpoints and such lines are aligned prior to ICP minimization. This whole process is very fast and reduces the complexity of break curve matching drastically by selecting corner candidates and their associated appropriate arcs for ICP alignment.

Since data from an excavation site has noise on a boundary, matching with *Horoscope* returns false alarms. True matches are founded among them if the match is coherent with the APC estimated by the previous step (Section 4).

6. Experiments

Three types of data are collected and experiments are performed for evaluating the system. First, synthetic data which have available ground truth is generated. Axially symmetric 3D point sets are generated by rotating 6 polynomial plane curves about an axis. Points are sampled at average distance 0.2mm and zero-mean Gaussian noise with variance 0.02 is added. The height of the synthetic data is between 200mm and 300mm along the axis, similar to a real pot. Then, each 3D data set is randomly broken into N fragments (data1). Second, to generate data from a real pot, a real pot is broken into 13 pieces and the outer surface of 10 fragments are densely scanned (data2). Third, 38 fragments are obtained from the archaeological site *Tell es-Safi* without any prior knowledge of the fragments. Outer surfaces are scanned by dense data laser scanner (data3). Mesh information is reconstructed by ball pivoting algorithm.

The reason why the system is reliable is the reasonably high accuracy of fragment axis of symmetry estimation. The axis estimation algorithm is evaluated with data 1, 2 and 3. Figure 7 shows an axis estimation result of data 2 and 3 and its sufficiency and distinctiveness. Figure 6 shows the axis estimation accuracy based on the fragments from 4 synthetic pots (data1). If the direction of axis estimation for a fragment is within 5 degrees of the ground truth, the algorithm scored as having made a successful estimate. As the number of fragments for a pot increases, the fragment size decreases and the axis estimation accuracy decreases. Figure 6 also shows that the axis estimation is more accurate after pairwise re-estimation.

The overall system is tested with data1 and its ground truth by increasing the number of broken fragments from a single pot (see table 1). As the number of fragments increases, each piece becomes smaller and hence the number of correct configurations estimated decreases.

One of the most notable benefits of the system is the capability of simultaneously configuring *a priori* unknown number of pots from fragments which are mixed together in an unorganized pile. Each pot may have missing an unknown number of its fragments. In addition, the system is designed to configure pots using only noisy outer surface measurement of real fragments. To measure algorithm performance, the system is tested with data2 and data3 ran-

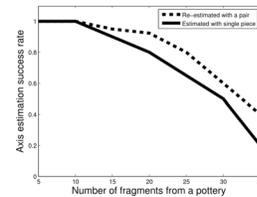


Figure 6: Axis estimation accuracy.

# of frag. from a pot	10	17	23	30
Correct conf.#	10/10	17/17	22/23	25/30

Table 1: Number of correct configurations as the number of fragments from a pot increases(data1).

	Break curve based	Axis based	Combin.
Conf.#	4/48	24/48	31/48

Table 2: Number of configurations for each steps alone and their combination(data2 and 3).

domly piled up together. The total number of fragments is 48. 26 pieces are classified as having distinctive APC property. Others are non-informative (Figure 7). 24 pieces are configured based on axis of symmetry, and 7 with break curve matching. The system automatically configures 3 pots out of the the 48 fragments. The three pots consist of 10, 16, 5 pieces respectively (see table 2). Figure 8a shows pairwise matching results based on axis of symmetry, while Figure 8b displays break curve matches. The straight lines are the common axes of symmetry. Figure 9 shows the final estimated results of 3 configurations. The system spends 634 minutes for the final configuration with Intel Core i7. ¹

7. Conclusion

Axially symmetric pots are automatically configured from their fragments by using two strong cues, surface symmetry about an axis and break curve information. A novel approach is presented to automatically match for the first time more challenging types of fragments. These fragment types include not only those with significant surface curvature but also those which are almost flat, and include those which are chipped and are represented by very noisy data. Our algorithm is also considerably faster than previous approaches in handling all types of fragments. It makes optimal use of geometric constraints to reduce fragment alignment complexity and introduces new insights to reduce the number of fragment candidate matches based on a new measure of break curve geometry. The estimation of multiple pots given a pile of all the fragments and the reasoning provided about their joint and pairwise geometries proved to be a powerful tool to deal with this puzzle problem. Experiments show that the system is robust to noise, bumps and erosion and estimates all pots axis/profile curves for which there is extractable information in the set of fragments.

¹All configurations from data 1, 2 and 3 can be found at http://www.lems.brown.edu/~sonkilho/Index_files/sup.pdf



Figure 7: Blue (left) and red (right) numbers are respectively sufficiency and distinctiveness for the axis estimation. Small blue values mean high sufficiency and high red values mean high distinctiveness. If the distinctiveness is bigger than 3.0, the system assume that the fragment has enough information for axis estimation. Blue points are center of circles fit to points within thick planar layers which are perpendicular to axis. Red points are inliers for estimating axis. The black line is the estimated axis.

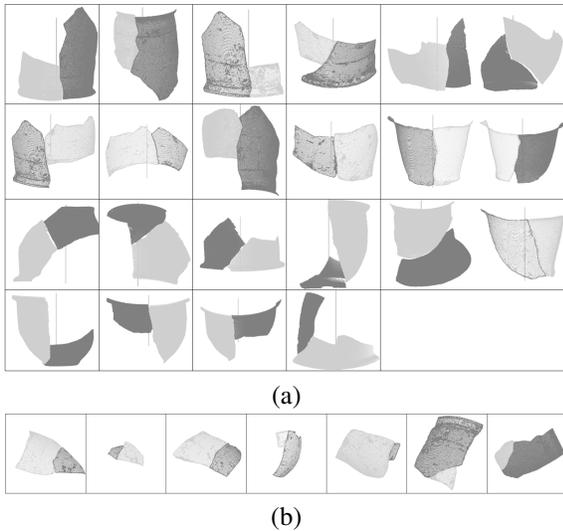


Figure 8: (a) Final pairs with *axis* based matching method, (b) Final pairs with *break curve* based matching method.

References

- [1] T. S. Cho, S. Avidan, and W. T. Freeman. A probabilistic image jigsaw puzzle solver, 2010. *Proc. CVPR*.
- [2] H. Freeman and L. Garder. Apictorial jigsaw puzzles: the computer solution of a problem in pattern recognition. *Electronic Computers*, 13:118–127, 1964.

- [3] A. Gallagher. Jigsaw puzzles with pieces of unknown orientation. In *Proc. CVPR*, 2012.
- [4] D. Goldberg, C. Malon, and M. Bern. A global approach to automatic solution of jigsaw puzzles. *Computational Geometry*, 28:165–174, June 2004.
- [5] D. Han, D. Cooper, and H. soo Hahn. Fast axis estimation from a segment of rotationally symmetric object. In *Proc. CVPR*, pages 1154–1161, June 2012.
- [6] Q.-X. Huang, S. Flöry, N. Gelfand, M. Hofer, and H. Pottmann. Re-assembling fractured objects by geometric matching. *ACM Trans. Graphics*, 25(3):569–578, 2006.
- [7] M. Kampel and R. Sablatnig. On 3d mosaicing of rotationally symmetric ceramic fragments. In *Proc. ICPR*, volume 02, pages 265–268, 2004.
- [8] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints, 2005. International Conference of Computer Vision (ICCV).
- [9] H. Mara and R. Sablatnig. Orientation of fragments of rotationally symmetrical 3d-shapes for archaeological documentation, 2006. 3DPVT.
- [10] J. C. McBride and B. B. Kimia. Archaeological fragment reconstruction using curve-matching. In *CVPR Workshop*, June 2003.
- [11] G. Papaioannou, E.-A. Karabassi, and T. Theoharis. Reconstruction of three-dimensional objects through matching of their parts. *PAMI*, 24(1):114–124, 2002.
- [12] D. Pomeranz, M. Shemesh, and O. Ben-Shahar. A fully automated greedy square jigsaw puzzle solver, 2011. *Proc. CVPR*.
- [13] A. Willis, D. Cooper, et al. Accurately estimating sherd 3D surface geometry with application to pot reconstruction. In *CVPR Workshop ACVA*, June 2003.
- [14] A. Willis and D. B. Cooper. Bayesian assembly of 3d axially symmetric shapes from fragments, 2004. *Proc. CVPR*.
- [15] A. Willis and D. B. Cooper. Estimating a-priori unknown 3d axially symmetric surfaces from noisy measurements of their fragments, 2006. 3DPVT.
- [16] A. Willis and D. B. Cooper. Computational reconstruction of ancient artifacts. *IEEE Signal Processing Magazine*, pages 65–83, 2008.
- [17] L. Zhu, Z. Zhou, and D. Hu. Globally consistent reconstruction of ripped-up documents. *PAMI*, 30(1):1–13, 2008.

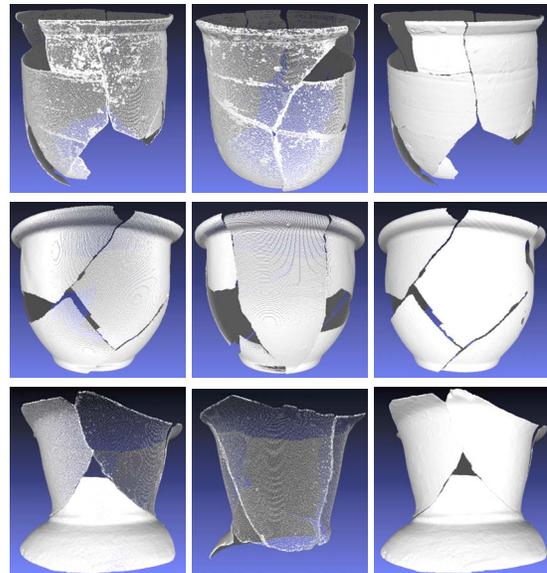


Figure 9: Three configurations out of 48 total fragments (16, 10, 5 pieces for each configuration).