

A Flexible Approach to Reassembling Thin Artifacts of Unknown Geometry

Geoffrey Oxholm (PhD Candidate)* and Ko Nishino (Associate Professor)

{geoff.oxholm, ko.nishino}@drexel.edu

Drexel University

Department of Computer Science

3141 Chestnut Street

Philadelphia, PA 19104

PHONE: +1-215-895-2669

FAX: +1-215-895-0545

A Flexible Approach to Reassembling Thin Artifacts of Unknown Geometry

Abstract

We present a novel 3D reassembly method for fragmented, thin objects with unknown geometry. Unlike past methods, we do not make any restrictive assumptions about the overall shape of the object, or its painted texture. Our key observation is that regardless of the object's shape, matching fragments will have similar geometry and photometry along and across their adjoining regions. We begin by encoding the scale variability of each fragment's boundary contour in a multi-channel, 2D image representation. Using this multi-channel boundary contour representation, we identify matching sub-contours via 2D partial image registration. We then align the fragments by minimizing the distance between their adjoining regions while simultaneously ensuring geometric continuity across them. The configuration of the fragments as they are incrementally matched and aligned form a graph structure that we use to improve subsequent matches. By detecting cycles in this graph, we identify subsets of fragments with interdependent alignments. We then minimize the error within the subsets to achieve a globally optimal alignment. We leverage user feedback to cull the otherwise exponential search space; after each new match is found and aligned, it is presented to a user for confirmation or rejection. Using ceramic pottery as the driving example, we demonstrate the accuracy and efficiency of our method on six real-world datasets.

1
2
3
4
5
6
7
8
9 **Research Aims**

10
11
12 Thin objects, such as hollow ceramic sculptures, bowls and vases, play an important role in
13 the study of cultural heritage. Unfortunately, reassembling such objects from their fragments
14 is challenging because the exposed “break-surfaces” are so thin. There are two general strate-
15 gies used to account for this missing information. *Global* approaches assume that the overall
16 shape of the object is known a priori; this greatly simplifies the process of aligning each piece.
17 *Local* approaches, rely on accurately aligning only two or three pieces at a time, building the
18 object up incrementally. Local methods are more general, since they do not rely on restrictive
19 assumptions about the object’s overall shape, but suffer from error accumulation. The goal of
20 this work is to reassemble thin objects without assuming a known global model and without
21 suffering from accumulated error. To do so we present a three step method in which accurate
22 local alignments are first estimated and then continually refined as the full model is uncovered.
23
24
25
26
27
28
29
30
31
32

33
34 **1. Introduction**
35
36

37 Cultural heritage objects, which are frequently broken when they are discovered, must be
38 reassembled before they may be interpreted. This process is so excessively time consuming
39 that many objects are never reassembled. The goal of virtual reassembly is to not only ease the
40 process of reassembling heritage objects, but also to enable wider access to these objects.
41
42
43
44

45 For thick objects, such as thick tile work and sculptures, past methods [1, 2, 3] rely on
46 the “break-surfaces” that have been exposed by the fractures. This information is valuable
47 because matching fragments of such objects will have oppositely shaped break-surfaces. Sets
48 of geometric features in these regions (such as sharp points, or valleys) may be compared to
49 determine if two fragments match each other. Once two fragments are determined to match, the
50 entire set of break-surface points can then be used to accurately align the fragments. For thin
51 objects, both the matching and aligning problems become much more challenging. Instead of
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 sharing oppositely shaped break-surfaces, fragments of thin objects share only a minute strip
10 in common.

11
12 This thin strip, which is often relatively featureless, is seldom descriptive enough to accu-
13 rately determine which fragments align with each other. In addition, since this matching region
14 is only a curve, and not a surface, it is insufficient in itself to correctly align two matching frag-
15 ments. In order to simplify these problems, past methods have relied on restrictive assumptions
16 about the geometry or painted texture of the object.
17
18
19
20
21

22 In this paper, we present our three-step method that reassembles objects using only the
23 fragments' boundary contours with minimal user interaction. Our method exploits the key ob-
24 servation that regardless of the shape, or painted texture of the object, the matching boundary
25 regions of adjoining fragments will be similar, both in geometry and photometry. Figure 1
26 outlines our approach. (a) First, we preprocess each fragment to encode the scale variability
27 of its boundary contour as a multi-channel 2D image. (b) We then identify matching sub-
28 contours using a novel image registration method based on these scale-space boundary contour
29 representations. (c) Next, we estimate the transformation to align the fragments using a least
30 squares formulation that minimizes the distance between the adjoining regions while simulta-
31 neously maximizing the resulting geometric continuity across them. (d) The configuration of
32 the fragments as they are incrementally matched and aligned form a graph structure. By iden-
33 tifying cycles in this graph, we detect subsets of fragments whose alignments are dependent on
34 each other. When a cycle is formed, we jointly re-optimize the alignments of the constituent
35 fragments to ensure a globally optimal configuration, and improve subsequent matches.
36
37
38
39
40
41
42
43
44
45
46
47
48

49 We use ceramic pottery as the driving example, and validate our method on several recently
50 excavated real-world historic artifacts. The results show that our method allows for accurate
51 reassemblies to be achieved with minimal human interaction, even when many fragments are
52 missing.
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 **2. Related Work**

10
11
12 When reassembling an object from its fragmented parts, the basic approach is to iterate
13 between finding matching fragments and bringing them into alignment. Reassembling thin
14 objects is particularly challenging because matching fragments share only their boundaries
15 in common. In other words, their “break-surfaces” are not surfaces at all, but contours. On
16 their own, these contours provide insufficient information to reliably match candidate frag-
17 ments since many fragments have similar shapes. This lack of information also makes aligning
18 matching fragments challenging. We cannot simply place the two matching boundary contours
19 next to each other since this does not guarantee that the fragment surfaces themselves will
20 reform the original unbroken shape.
21
22

23
24
25 To address the ill-posed problems of finding matching fragments and aligning them, a com-
26 mon approach is to restrict the class of objects that may be reassembled by imposing a global
27 model. Some authors [4, 5, 6] assume the object is axially symmetric. The shape of the frag-
28 ment can then be compared to the profile contour of the object model to help determine its
29 proper location and orientation. This effectively reduces the problem to a sort of curved 2D
30 puzzle. To determine the profile contour of a vase or bowl, the authors also assume that the
31 fragments themselves are able to reliably convey this information. If the fragments are too
32 small, the profile contour may not be estimated. We avoid relying on simple object models,
33 or assuming sufficiently large fragment sizes by exploiting the geometric continuity that must
34 exist across matching fragments.
35
36

37
38
39 Even when a global model is used to simplify the reassembly process, methods that rely on
40 pairwise alignments necessarily suffer from error accumulation. This effect can even be seen
41 when the object has only three pieces. Once the first two pieces have been matched, the space
42 for the third piece may already be too small or too large. To address this, some authors [7, 5]
43 delay the alignment phase until clusters of three matching fragments have been found. While
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 this does improve alignment accuracy, it is still a local optimization and therefore prone to error
10 accumulation.
11

12
13 In range image registration, where pairwise alignment is much more reliable, accumulated
14 error may be evenly dispersed to finalize a reassembly. Once a series of overlapping fragments
15 is found that connects back to itself, Sharp et al. [8] then divide up the resulting gap or over-
16 lap over the whole set of fragments. Pulli [9] uses the point correspondences of the pairwise
17 alignments as soft constraints on a final global alignment. In object reassembly, however, the
18 pairwise alignments are not as reliable. In addition to improving the overall reassembly accu-
19 racy, each incremental match provides important contextual clues that can be used to improve
20 subsequent matches. By leveraging the context provided by past matches, we determine the
21 placement of relatively nondescript fragments, and ensure global accuracy.
22
23

24
25 Sađirođlu and Erđil [10] focus on the fragments' painted texture. If the object is sufficient
26 patterned, they use texture synthesis and inpainting methods to expand each fragment's surface
27 texture. This can be thought of as generating overlapping surfaces which may then be aligned
28 using standard texture matching approaches. We also use the texture painted onto the surface,
29 but only as part of our representation of the boundary contour and do not assume any a priori
30 knowledge of the pattern or characteristics of the texture.
31

32
33 Some authors focus specifically on the problem of finding matching boundary contours of
34 thin fragments using only their 3D geometry. Although the contours are three dimensional,
35 Wolfson [11] showed that finding matching sub-contours can be recast as a one-dimensional
36 string matching problem. This dimensionality reduction is possible because contours can be
37 uniquely expressed in terms of their rotationally and translationally invariant geometric char-
38 acteristics: curvature and torsion. By encoding contours in this way, matching regions may be
39 detected as matching series of curvature and torsion tuples using the longest common substring
40 algorithm.
41
42

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Unfortunately, recasting 3D contour matching as 1D string matching is highly sensitive to noise, chipping and other realities since it relies on the calculation of the contour’s third derivative. To address this, including scale in the matching process has been quite successful. Large scale detail can be used to estimate matches, while finer scale detail can be used to refine and validate the match. The standard approach to incorporate scale is to incrementally smooth the 3D points of the contour, recalculating the curvature and torsion values at each degree of smoothness [11, 12, 13, 14, 15]. Although incremental smoothing highlights features of prominence and mitigates the effect of noise, smoothing the *geometry* of the contour introduces shapes that are dramatically different from the original. Our approach is more faithful to the notion of scale. By smoothing the geometric *characteristics* of the contour we induce an effect similar to moving away from the object. In addition to this, we incorporate the photometric characteristics of the boundary contour. This additional information helps inform the contour matching process, and is particularly useful when many boundary contours share similar geometry.

3. Boundary Contour Representation

The first step in our method is to identify which fragments are most likely to align, and where their boundaries match. To do so quickly and accurately, we leverage the scale variability of each fragment’s boundary contour. A coarse scale representation, which is robust to noise and subtle detail, may be used to quickly estimate potential matches, while finer scale detail may be used to verify and fine-tune the estimated matches. This graduated relationship naturally lends itself to a hierarchical encoding. To that end, we build a multi-channel image representation that encodes the scale variability of each fragment’s shape and color.

3.1. Boundary Extraction

As shown in Figure 2, we start by acquiring geometric and photometric information about each fragment using a range sensor; specifically a light-stripe range scanner Canon VIVID 910. Without loss of generality, we assume the entire surface of each fragment may be viewed from a single viewpoint. Although this assumption holds true for the various objects in our datasets, if it were not the case, multiple scans could be used to form the exterior surface of the fragment. Each scan is comprised of a 640×480 color image (Figure 2a), as well as three-dimensional coordinates for each of the image points (represented in Figure 2b as a depth map). In most cases we use depth and color discontinuities to automatically isolate the fragment in the scan. The result is a 2D mask like the one shown in Figure 2c. Occasionally, part of the thin break-surface is visible and manual intervention is required to correct the mask.

Using Suzuki’s border following method [16], we extract the 2D border of the mask (shown in Figure 2d). We then extract the 3D location of each point along the border from the range data. This contour (shown in Figure 2e) is noisy and its sampling is too dense to allow for reliable analysis. To address this, we smooth the contour slightly and sub-sample the points using an iterative technique described by Leitao and Stolfi [13] which ensures that the points along the resulting contour are uniformly spaced. The processed boundary contour (Figure 2f) exhibits a necessary degree of smoothness to reliably evaluate while remaining true to the original shape. Once the geometry of the 3D contour has been processed, we project the contour back to the image plane. Here we extract the color of each contour point using bilinear interpolation.

3.2. Boundary Representation

We describe each fragment’s boundary contour as a cyclic string $f(t)$ of four-valued feature vectors

$$f(t) = \{(\kappa, \tau, c_r, c_g)_{t \bmod n}\}, \quad (1)$$

where n is the number of samples along the contour and t is the sample index. The first two values, κ and τ , are curvature and torsion. These two values encode the rotationally and translationally invariant geometry of the 3D contour. The second two, c_r and c_g , are the red and green chromaticity; they encode the appearance of the contour. These values help locate and refine matches, particularly when the contour geometry is relatively featureless. We use chromaticity because it provides a measure of invariance to illumination conditions, which is an important consideration in many domains. Also, observe that the string $f(\cdot)$ is periodic, i.e., $f(t+n) = f(t)$. This captures the cyclic nature of the contour.

The curvature and torsion (κ and τ) of a 3D contour $\lambda(t) = (x(t), y(t), z(t))$ can be defined in terms of the derivatives of the three coordinates [17] as,

$$\kappa = \frac{\|\ddot{\lambda} \times \dot{\lambda}\|}{\|\dot{\lambda}\|^3} = \frac{\sqrt{A^2 + B^2 + C^2}}{(\dot{x}^2 + \dot{y}^2 + \dot{z}^2)^{3/2}} \quad (2)$$

and

$$\tau = \frac{(\dot{\lambda} \times \ddot{\lambda}) \cdot \ddot{\lambda}}{\|\dot{\lambda} \times \ddot{\lambda}\|^2} = \frac{\begin{vmatrix} \dot{x} & \dot{y} & \dot{z} \\ \ddot{x} & \ddot{y} & \ddot{z} \\ \ddot{x} & \ddot{y} & \ddot{z} \end{vmatrix}}{A^2 + B^2 + C^2}, \quad (3)$$

where

$$A = \begin{vmatrix} \dot{y} & \dot{z} \\ \ddot{y} & \ddot{z} \end{vmatrix}, \quad B = \begin{vmatrix} \dot{z} & \dot{x} \\ \ddot{z} & \ddot{x} \end{vmatrix}, \quad C = \begin{vmatrix} \dot{x} & \dot{y} \\ \ddot{x} & \ddot{y} \end{vmatrix}. \quad (4)$$

Here $\|\square\|$ and $|\square|$ denote the $L2$ norm and matrix determinant, respectively. $\dot{\square}$, $\ddot{\square}$, and $\dddot{\square}$ denote the first-, second-, and third-order derivatives with respect to the arc length t , i.e., $\frac{\partial}{\partial t}$, $\frac{\partial}{\partial t^2}$ and $\frac{\partial}{\partial t^3}$. We compute these using central numerical differentiation.

These geometric values are used frequently in shape description applications because they provide a unique, rotationally and translationally invariant representation of the contour. We

1
2
3
4
5
6
7
8
9 extend the descriptive formulation beyond past work by additionally encoding the photometric
10 properties of the contour. After white-balancing the color image, we then encode the color of
11 each contour point using its red and green chromaticity values (c_r and c_g respectively). These
12 are computed by simply dividing the color channels by the total intensity of the point. Note
13 that including the blue channel would be redundant since the three values sum to one.
14
15
16
17

18 3.3. Encoding Scale

19
20 Although this string of values $f(t)$ accurately describes the boundary contour, it encodes
21 small-scale detail and noise that may not precisely align with the matching sub-contour de-
22 scription of an adjoining fragment. In order to reliably locate matching boundary sub-contours,
23 we exploit the scale variability of their geometry and photometry. We use the coarse scale rep-
24 resentation, which is robust to noise and subtle detail, to quickly estimate potential matches,
25 and the finer scale detail to verify and fine-tune them.
26
27
28
29
30
31
32

33 Past authors [11, 12, 13, 14, 15] typically encode scale by incrementally smoothing the
34 geometry of the contours themselves. This approach, however, introduces shapes that are dra-
35 matically different from the original. In order to maintain the authenticity of the underlying
36 geometry and photometry of the contour our approach is to iteratively smooth the string $f(t)$
37 itself. We then store the smoothed values as rows of a multi-channel image \mathbf{S} , like the one
38 shown in Figure 3. For clarity, we have shown each channel separately. Left-to-right, top-to-
39 bottom they are curvature, torsion, red and green chromaticity. The lowest row of pixels \mathbf{S}_0
40 encodes the smallest scale detail, i.e., $\mathbf{S}_0 = f(t)$. Moving up the image to row r corresponds to
41 a coarser scale. Each increasing scale is calculated using circular convolution of the base row
42 with a Gaussian smoothing kernel $\mathcal{N}(\sigma_r)$ of standard deviation σ_r proportional to r ,
43
44
45
46
47
48
49
50
51
52

$$53 \mathbf{S}_r = \mathcal{N}(\sigma_r) \circledast \mathbf{S}_0, \quad (5)$$

54 where \circledast is the circular convolution operator (convolution with periodic a boundary). In our
55
56
57
58

1
2
3
4
5
6
7
8
9 case, we let $\sigma_r = 0.05r$ and build images with 100 rows. Increasing the number of rows,
10 and consequently the maximum value for σ , will increase the number of false positives as
11 discerning detail becomes smoothed away. Decreasing the number of rows, on the other hand,
12 will lead to false negatives and ultimately, incomplete reassemblies. We therefore chose this
13 value empirically to minimize the false positives without allowing for many false negatives.
14
15
16
17

18 The resulting representation is now a 4-channel 2D image that encodes the scale variability
19 of the boundary contour's geometry and photometry. In Figure 3, observe how the five sharp
20 corners of the fragment (e) have become bright regions in the curvature channel (a), and how
21 the banded painting yields a similarly striped pattern in the red chromaticity channel (c). Note
22 that as torsion (b) is a signed value, in this channel, a value of 0 is given an intensity of 0.5.
23 Although the image has four channels, we use a more compact rendering throughout this paper.
24 As shown in (f), we replace the two chromaticity channels with a single channel computed as
25 the intensity. This channel is colored blue; curvature is shown in red, and torsion is green.
26
27
28
29
30
31
32
33

34 This 2D formulation is particularly appealing because it describes the photometric and geo-
35 metric properties of the contour under all scales at once. It is this characteristic that we leverage
36 when solving the matching problem.
37
38
39
40

41 **4. Matching Boundary Contours**

42
43

44 By encoding the scale variability of each boundary contour's shape and color as a multi-
45 channel image, the problem of finding matching sub-contours is now akin to partial image
46 registration. Specifically, matching image regions will correspond to matching sub-contours.
47 Note that because both contours are encoded in counterclockwise ordering, one image must
48 first be horizontally flipped. Although chromaticity has a standard range of $[0, 1]$, curvature
49 and torsion are unbounded. To introduce a degree of comparability across all channels, we
50 limit each range by scaling it. Figure 4 outlines our three step matching method.
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 **a) Longest common sub-contour** Similar to past work [14, 15, 7], we begin our search with
10 the coarsest scale, i.e., the top row. Here we perform a longest common substring analysis with
11 two modifications. First, a pair of four-valued feature vectors $r_0^A(p)$ and $r_0^B(q)$ is considered to
12 be matching if each of their values differ by less than a threshold. Secondly, we allow matches
13 to wrap around the images. The table produced by the dynamic programming algorithm is then
14 traversed, and matches are inserted into a queue where priority is given based on the length of
15 the match. Figure 4a, shows the top match for the two fragments shown.
16
17
18
19
20
21

22
23 **b) 2D image registration refinement** Since this coarse scale involves the most smoothing,
24 there remains some ambiguity about the precise location of the matching sub-contour. Just
25 as archaeologists do by hand, we leverage finer-scale detail to validate the match. Using
26 one region as a template, we refine the location of the matching region via normalized cross-
27 correlation image registration. In Figure 4b we show the result of this step which resembles
28 the “lock-in” effect archaeologists use to describe the certainty that comes from aligning small
29 scale features. Note that the matching region of the smaller piece has moved right in the scale-
30 space representation and clockwise around the piece.
31
32
33
34
35
36
37

38
39 **c) Fan-out estimation** The first and last points of every correctly matching contour sit next
40 to a non-matching point. In Figure 4, for example, the match stops when the two contours turn
41 sharply away from each other. This results in contrasting values in the two contour descrip-
42 tion strings encoded at the base of the scale-space representations. As we increase the scale,
43 and move up the scale-space representations, these non-matching values are smoothed with
44 an increasing number of matching values (as the standard deviation of the smoothing function
45 σ_r increases). The result of this smoothing is an iterative decreasing in the number of match-
46 ing values. To address this effect, which we call *fan-out*, we incrementally expand the base
47 of the matching regions until a threshold is surpassed. As shown in Figure 4c, the resulting
48 trapezoidal regions correctly convey the full match.
49
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 **5. Pairwise Alignment**

10
11
12 Once two matching boundary regions have been identified, the second step of our algorithm
13 is to estimate the transformation that brings the fragments into alignment. We formulate this as
14 a least-squares optimization problem where the error is measured at each sampled point of the
15 matching boundary contours according to two metrics. The first, which we call “contour error,”
16 quantifies the distance between the corresponding points. This ensures that the fragments are
17 tightly aligned. The second, which we call “surface error,” quantifies the geometric continuity
18 across the points. This ensures that the resulting surface geometry transitions smoothly from
19 one fragment to the other.
20
21
22
23
24
25

26
27 More precisely, given two fragments \mathbf{A} and \mathbf{B} , we seek the transformation $T_{\mathbf{B},\mathbf{A}}$ that brings
28 \mathbf{B} into alignment with \mathbf{A} such that the sum of squared residuals $\sum_i^m e_i^2$ across the m points
29 of their matching boundary contours is minimized. We formulate the residual at point i as a
30 weighted combination of our two metrics,
31
32
33

34
35
36
$$e_i = e_i^c + \alpha e_i^s, \tag{6}$$

37
38 where e_i^c and e_i^s are the contour error and surface error, respectively, and α is a relative weight.
39
40

41 *Contour error*

42
43 To formulate the contour error e^c , we let $\{\mathbf{a}_i \mid 1 \leq i \leq m\}$ be the m 3D contour points
44 from \mathbf{A} , and let $\{\mathbf{b}_i \mid 1 \leq i \leq m\}$ be the corresponding points from \mathbf{B} . The contour error e_i^c at
45 point i is then the Euclidean distance between corresponding points,
46
47
48
49

50
51
$$e_i^c = \|T_{\mathbf{B},\mathbf{A}}(\mathbf{b}_i) - \mathbf{a}_i\|. \tag{7}$$

52
53 The transformation $T_{\mathbf{B},\mathbf{A}}$ is defined in terms of a rotation matrix \mathbf{R} and a translation vector \mathbf{t} ,
54
55

56
$$T_{\mathbf{B},\mathbf{A}}(\mathbf{b}_i) = \mathbf{b}_i\mathbf{R} + \mathbf{t}. \tag{8}$$

This error metric is insufficient on its own because the contours do not convey any reliable information about the fragment surfaces themselves; the alignment should not only be tight, but should result in a smooth surface.

Surface error

Our second error term evaluates the alignment quality of the fragment surfaces. We evaluate the continuity of the newly-formed surface in terms of the progression of surface-normal vectors across the matching boundary contour points. In Figure 5 we illustrate how the surface error e_i^s is computed for a point i . We begin by computing the plane P_i that is orthogonal to the boundary contour’s tangent vector at this point. The intersection of this plane with both surfaces forms a curve (shown in red and blue). Along this curve, we extract surface normals from the fragment surfaces at regular intervals. Specifically, we let $\hat{\mathbf{a}}_i(j)$ denote the j^{th} surface normal from point \mathbf{a}_i , and define $\hat{\mathbf{b}}_i(j)$ analogously. We then form a piecewise function $\mathbf{q}_i(k)$ of these normal vectors as

$$\mathbf{q}_i(k) = \begin{cases} \hat{\mathbf{a}}_i(-k) & \text{for } k < 0 \\ \mathbf{0} & \text{for } k = 0 \\ T_{\mathbf{B},\mathbf{A}}(\hat{\mathbf{b}}_i(k)) & \text{for } k > 0, \end{cases} \quad (9)$$

where $T_{\mathbf{B},\mathbf{A}}$ is the alignment transformation to be optimized. Note that we negate k in the first case since $\hat{\mathbf{a}}(\cdot)$ is only defined for positive inputs.

An optimal alignment will result in a gradient of surface normals $\frac{\partial \mathbf{q}_i}{\partial k}$ whose value at the boundary (i.e., $k = 0$) matches the gradient of surface normals on the adjoining fragments. We therefore compute a target gradient value equal to the mean of the gradients from both fragments evaluated a short distance ϵ from the boundary contour. The surface alignment residual

error is then computed as

$$e_i^s = \frac{\partial \mathbf{q}_i}{\partial k}(0) - \frac{1}{2} \left[\frac{\partial \widehat{\mathbf{a}}_i}{\partial j}(\epsilon) + \frac{\partial \widehat{\mathbf{b}}_i}{\partial j}(\epsilon) \right]. \quad (10)$$

To compute the derivative of \mathbf{q} at point i , we extract fifteen normals $\{\mathbf{q}_i(k) \mid -7 \leq k \leq 7\}$ from a thin strip of the corresponding surfaces. By restricting this calculation to a thin strip we avoid any assumptions about the global geometry of the object. We then weight these 15 normals using a standard discrete derivative kernel. For the target gradients $\partial \widehat{\mathbf{a}}_i$ and $\partial \widehat{\mathbf{b}}_i$, we use the same process and evaluate the gradient at the midway point $\epsilon = 4$. We then solve this system with Levenberg-Marquardt iterative minimization [18, 19].

To provide a measure of similarity between the two error terms e^c and e^s , we set the relative weighting term α (from Equation 6) equal to the boundary contours' sub-sampling distance divided by π (the maximum value for e^c). To reduce the running time of this optimization, we pre-compute the intersection planes, surface normals, and target gradients for each contour point. By doing so, however, we are making the assumption that the intersection plane computed for fragment A at point i is parallel to the intersection plane computed for fragment B . This assumption may be violated, in particular when alignments are re-optimized, as we discuss next. To address this, we add an additional error term that quantifies the orientation error of these planes. Since the range of this error is the same as for e^s we also weight it with α .

6. Incorporating Context

As illustrated in Figure 6, the final component of our framework is to incorporate the global geometry of the object as it becomes known. We use each additional match to increase the overall alignment quality of the object, which in turn improves the quality of future matches. We encode the global geometry as a graph in which the nodes represent individual fragments, and the edges correspond to matching contours between them. Cycles in this graph correspond

1
2
3
4
5
6
7
8
9 to subsets of fragments whose alignments depend on each other, i.e. changing the orientation
10 of one fragment in such a subset will impact the alignment quality of the others.
11

12
13 Figure 6 shows the significance of a cycle in the graph. In the four node chain of Figure 6a
14 there is a gap between fragments *B* and *C* that has resulted from the accumulation of small
15 alignment errors in the three edges. In Figure 6b, the addition of the edge connecting fragments
16 *B* and *C* has created a cycle. It is plain to see that aligning the two fragments using only this
17 newly found matching contour (marked in yellow) will decrease the quality of the neighboring
18 alignments. It would also be inaccurate to simply move fragment *B* evenly between its neigh-
19 bors. Instead, this additional edge provides important context information that should be used
20 to improve all of the other alignments.
21
22

23
24 We detect cycles by computing all of the paths between the fragments involved in the new
25 match. This is done simply with a single depth first traversal of the graph. Using the corre-
26 spondences embedded in every edge of the cycle to evaluate the overall error in the subset, we
27 jointly re-optimize the constituent alignments. Note how the gap between *B* and *C* has been
28 closed by this process, and the overall alignment quality of the subset has been improved.
29

30
31 We formulate the update alignment of the subset as a least squares optimization where
32 the residual error is measured across all points of every matched contour pair according to
33 Equation 6. All alignment translations and rotations are optimized jointly to minimize the error
34 of the entire subset. As the number of cyclically dependent nodes approaches the entire set of
35 fragments, this re-optimization becomes a full global optimization. By building up the global
36 model in a bottom-up fashion, however, each re-optimization is simply a refinement of the
37 previous configuration. Thus, we leverage the accurate estimates of our pairwise alignments to
38 reduce the search-space of the global optimization.
39
40

41
42 After updating the alignment of the subset, the set of points being used to evaluate an align-
43 ment may no longer be an optimal set of correspondences. Recall that the alignment quality
44
45
46
47
48
49
50
51
52

1
2
3
4
5
6
7
8
9 across a pair of matching contours depends on the distance between pairs of corresponding
10 points. When a cycle is formed, we may find that one fragment should slide along the other to
11 improve the overall alignment quality. To that end we utilize the core functionality of Besl and
12 McKay’s iterative closest point (ICP) algorithm [20]. Specifically, we iterate between optimiz-
13 ing the transformations of the subset, and updating the correspondences used to evaluate them.
14 At each iteration every point is paired with the nearest point on the corresponding contour.
15 Once the correspondences become stable, the optimal alignment for the object will be found.
16
17
18
19
20
21
22
23

24 **7. Grouping Fragments**

25
26 The final aspect of our method is designed to aid in the detection of matching contour
27 groups such as the one shown in Figure 7. When two or more pieces are bound together
28 by matching sub-contours, they effectively form a single *meta-piece*. Intuitively, this corre-
29 sponds to virtually gluing pieces together—the matched sub-contours become internal to the
30 meta-piece. Each connected component subgraph of our alignment graph is given its own
31 scale-space boundary representation which is then compared against the remaining unpaired
32 fragments and meta-pieces.
33
34
35
36
37
38
39

40 The boundary representation for a meta-piece is formed by first extracting the unmatched
41 boundary points from all constituent pieces. These points are concatenated to form a single
42 boundary contour. As shown in Figure 7, there are two important factors to consider during this
43 process. First, the matching boundary contour between the two fragments may have been too
44 short. This will result in small sub-contours that essentially overlap each other as in Figure 7b.
45 To account for this, we extend the boundary contours as shown in Figure 7c. The second factor
46 to consider comes from our preprocessing step. Recall that each fragment’s raw boundary
47 contour is initially smoothed and sub-sampled, causing a slight rounding of sharp corners. As
48 shown in Figure 7c, the meeting point of two fragments must be smoothed to account for this.
49
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 This approach ensures that contour regions that have previously been matched are no longer
10 considered for future pairings. Additionally, the boundary contour of a meta-piece will typi-
11 cally have a more favorable representation. In Figure 7d, the third fragment is successfully
12 paired with both of the other fragments. Since we prioritize candidate matches based on their
13 length, the small matching contour that makes up the bottom of this match would not otherwise
14 have been found in a reasonable time-frame. Further examples can be seen in the final steps of
15 Figure 8, and in the beginning of the third row of Figure 12.
16
17
18
19
20
21
22
23

24 **8. Experimental Results**

25
26 To validate our algorithm we augmented it with an simple, interactive user interface. At
27 each iteration of our algorithm (summarized in Figure 1) we first extract the longest matching
28 boundary contour. We then align the constituent fragments and present the updated reassembly
29 to the user who may then accept or reject it. When a match is accepted, it is added to the re-
30 assembly, and the additional context information is propagated to the other alignments. Except
31 where indicated, the resulting reassemblies contain every fragment in the dataset. Each dataset
32 is incomplete, describing only a partial vessel. Despite this, we are able to successfully recover
33 the full 3D geometry of every dataset. Unfortunately, there is no ground truth to which we
34 could compare the exact alignment of the fragments.
35
36
37
38
39
40
41
42
43

44 In Figure 8 we show the partial reassembly of a store bought vase. At each step, the current
45 candidate match is shown in green, other matches that have been re-optimized are colored in
46 purple, and any other past matches are shown in red. In this artificial example, each of the
47 suggested matches is correct making the full reassembly essentially fully automatic. As such,
48 the full process requires only about 12 seconds.
49
50
51
52

53 In Figure 9 we summarize our results and compare total reassembly times using our method
54 with manual reassembly of the actual artifacts. In this graph, the store bought vase of Figure 8
55
56
57
58

1
2
3
4
5
6
7
8
9 referred to as VASE. The remaining five objects are recently excavated cultural heritage arti-
10 facts, and as such exhibit increased chipping, signs of wear, and a greater range of fragment
11 size. Although these datasets describe axially symmetric objects, the fragments themselves are
12 quite small. Consequently, past methods that rely on estimating the profile contour of symmet-
13 ric objects [4, 5] are not be well suited for this task.
14
15
16
17

18
19 In Figure 10 we show the steps taken by a user to reassemble 15 pieces of one such historical
20 artifact. Note, however, that we have omitted a small triangular piece in the top of the rim. This
21 omission is visible in the final reassembly (shown in a green box) as compared to the ground
22 truth image (shown in an orange box). This piece was too small to be reliably analyzed, due to
23 the fan-out effect discussed in Section 4. This example is referred to as BOWL1 in Figure 9.
24
25
26
27

28
29 In Figure 11 we show the reassembly of an 11 piece historical artifact. Note how a large
30 gap is closed in the last step, as indicated by the yellow circles. Since many pieces are missing
31 on this side of the bowl, a long chain of pieces has accumulated error. When the matching
32 contour is found that closes this gap, the entire rim of the vessel is re-aligned to distribute the
33 error. This example is referred to as BOWL2 in Figure 9.
34
35
36
37

38
39 Figure 12 shows our most challenging dataset. Whereas the painted scene provides use-
40 ful clues to archaeologists working with the artifact itself, the color information contained at
41 the boundary of each piece is too subtle to be discerning in our scale-space images. Further,
42 there are many small missing fragments that partition long matching contours into smaller
43 matches. Because of these two factors many erroneous matches are considered. We discard
44 clearly erroneous matches by inspecting the resulting alignment error. If too much space exists
45 between matching contours, or the resulting surface is not smooth, the match is discarded. This
46 dramatically reduces the amount of user interaction required in all cases. In this case, how-
47 ever, evaluating alignment error may involve large cycles of fragments. Note that almost every
48 step involves a large scale re-optimization. As shown in Figure 9, this vessel (referred to as
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9 PLATE1) is the only example where assembly by hand took less time than using our method.
10 Nevertheless, despite the challenges in this dataset, our method is able to rapidly align all but
11 one fragment of this object.
12
13

14
15 The final two datasets, which are respectively labeled PLATE2 and PLATE3 in Figure 9,
16 show the importance of our surface error alignment metric. Many of the matching contours
17 are essentially straight lines. As such, simply minimizing the distance between the contours
18 would result in dramatic surface discontinuity across the matching regions. By additionally
19 optimizing the geometric continuity across the surfaces, the pieces are accurately aligned.
20
21
22

23
24 Figure 9 shows the reassembly times for our six datasets using our method (light) and by
25 hand (dark). Although our user interface is rudimentary, it is important to note that with the
26 exception of PLATE1 our system allows for more rapid reassembly than by hand. BOWL2,
27 which is detailed in Figure 11, proved too challenging to reassemble by hand; the time reported
28 is when the test subject abandoned the task. We believe this is due to the object’s many similarly
29 shaped fragments. By utilizing the full scale-space of each fragment, we exploit subtle detail
30 to eliminate the guesswork associated with manual reassembly.
31
32
33
34
35
36
37
38

39 **9. Conclusion**

40
41
42 Our method, which makes no assumptions about the global geometric or photometric struc-
43 ture of the object, is applicable to any reassembly problem where the boundary of each fragment
44 may be extracted. This includes volumetric object reassembly, with the added benefit that a full
45 3D model of each fragment is not needed. To achieve this, we have presented a method that
46 balances our novel pairwise alignment method with a flexible global approach that leverages
47 the important context information each match provides. By allowing the global model to be
48 uncovered incrementally, we successfully complete full reassemblies even when most of the
49 object’s fragments are missing. Since access to cultural heritage objects is greatly restricted,
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 achieving reassembly of real artifacts in a virtual setting is itself an important step. For future
10 work, we plan to investigate more engaging and efficient user interface designs.
11
12

13 *Acknowledgements*

14
15 This work was supported in part by National Science Foundation CAREER Award IIS-0746717 and
16 IIS-0803670. The authors are grateful to Jed Levine of the National Park Service, and to Patrice Jeppson,
17
18 Glen Muschio and his students at Drexel University for their time and help with data acquisition.
19
20
21
22

23 **References**

- 24
25
26 [1] S. Winkelbach, F. M. Wahl, Pairwise Matching of 3D Fragments Using Cluster Trees, *International*
27 *Journal of Computer Vision* 78 (2008) 1–13.
28
29 [2] B. J. Brown, C. Toler-Franklin, D. Nehab, M. Burns, D. Dobkin, A. Vlachopoulos, C. Dumas,
30 S. Rusinkiewicz, T. Weyrich, A System for High-Volume Acquisition and Matching of Fresco
31 Fragments, *ACM Transactions on Graphics* 27 (2008) 1–9.
32
33 [3] Q. Huang, S. Flöry, N. Gelfand, M. Hofer, H. Pottmann, Reassembling Fractured Objects by
34 Geometric Matching, *ACM Transactions on Graphics* 25 (2006) 569–578.
35
36 [4] A. R. Willis, D. B. Cooper, Bayesian Assembly of 3D Axially Symmetric Shapes from Fragments,
37 *IEEE Conference on Computer Vision and Pattern Recognition* (2004) 82–89.
38
39 [5] The SHAPE Lab, Assembling Virtual Pots from 3D Measurements of Their Fragments, *Virtual*
40 *Reality, Archeology, and Cultural Heritage* (2001) 241–254.
41
42 [6] M. Kampel, R. Sablatnig, On 3D Mosaicing of Rotationally Symmetric Ceramic Fragments, *In-*
43 *ternational Conference on Pattern Recognition* 2 (2004) 265–268.
44
45 [7] K. Weixin, B. Kimia, On Solving 2D and 3D Puzzles Using Curve Matching, *IEEE Conference*
46 *on Computer Vision and Pattern Recognition* (2001) 583–590.
47
48 [8] G. C. Sharp, S. W. Lee, D. K. Wehe, Multiview Registration of 3D scenes by Minimizing Error
49 Between Coordinate Frames., *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26
50 (2004) 1037–50.
51
52
53
54
55
56
57
58

- 1
2
3
4
5
6
7
8
9 [9] K. Pulli, Multiview Registration for Large Data Sets, International Conference on 3-D Digital
10 Imaging and Modeling (1999) 160–168.
11
12 [10] M. Sağıroğlu, A. Erçil, A Texture Based Approach to Reconstruction of Archaeological Finds,
13 Symposium on Virtual Reality, Archaeology and Cultural Heritage (2005) 1–6.
14
15 [11] H. Wolfson, On Curve Matching, IEEE Transactions Pattern Analysis and Machine Intelligence
16 12 (1990) 483–489.
17
18 [12] E. Kishon, T. Hastie, H. Wolfson, 3-D Curve Matching Using Splines, European Conference on
19 Computer Vision 427 (1990) 589–591.
20
21 [13] H. Gama Leitao, J. Stolfi, A Multiscale Method for the Reassembly of Two-Dimensional Frag-
22 mented Objects, IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (2002) 1239–
23 1251.
24
25 [14] F. Mokhtarian, A. Mackworth, Scale-Based Description and Recognition of Planar Curves and
26 Two-Dimensional Shapes, IEEE Transactions Pattern Analysis and Machine Intelligence (1986)
27 34–43.
28
29 [15] F. Mokhtarian, Multi-scale Description of Space Curves and Three-Dimensional Objects, IEEE
30 Conference on Computer Vision and Pattern Recognition (1988) 298–303.
31
32 [16] S. Suzuki, K. Be, Topological Structural Analysis of Digitized Binary Images by Border Follow-
33 ing, Computer Vision, Graphics, and Image Processing 30 (1985) 32–46.
34
35 [17] D. Struik, Lectures on Classical Differential Geometry, Addison-Wesley, Cambridge, Mass., 1950.
36
37 [18] K. Levenberg, A Method for the Solution of Certain Non-linear Problems in Least-Squares, Quar-
38 terly of Applied Mathematics 2 (1944) 164–168.
39
40 [19] D. W. Marquardt, An Algorithm for Least-Squares Estimation of Nonlinear Parameters, SIAM
41 Journal on Applied Mathematics 11 (1963) 431–441.
42
43 [20] P. Besl, H. McKay, A Method for Registration of 3-D Shapes, IEEE Transactions on Pattern
44 Analysis and Machine Intelligence 14 (1992) 239–256.
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Captions for figures

1	System overview	3
2	Boundary contour extraction	7
3	Boundary contour representation	9
4	Boundary contour matching	12
5	Pairwise alignment	15
6	Incorporating context	
7	Grouping fragments	
8	VASE reassembly	
9	Reassembly times in minutes	
10	BOWL1 reassembly	
11	BOWL2 reassembly	
12	PLATE1 reassembly	

Figure 1

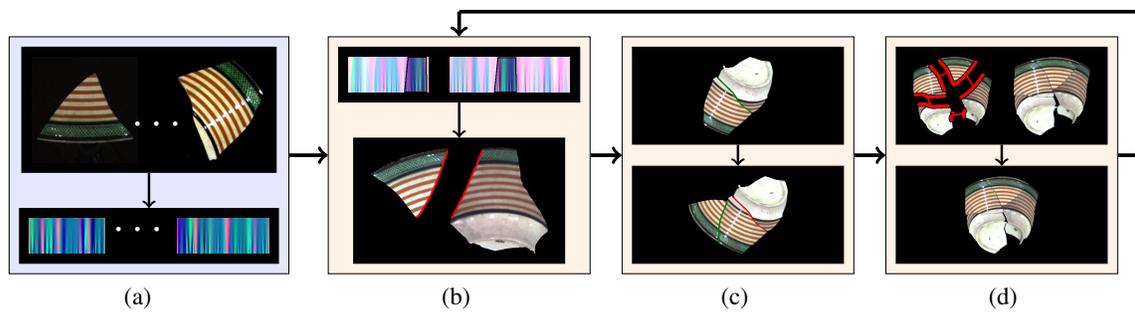


Figure 2

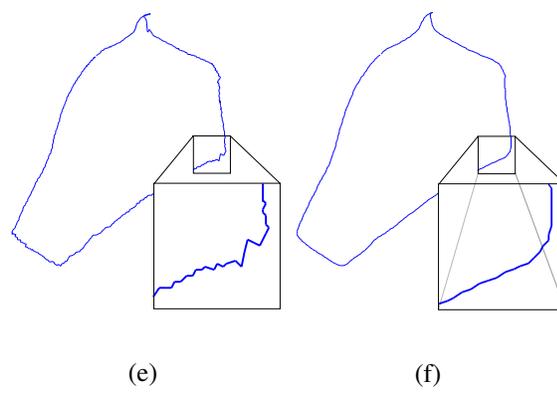
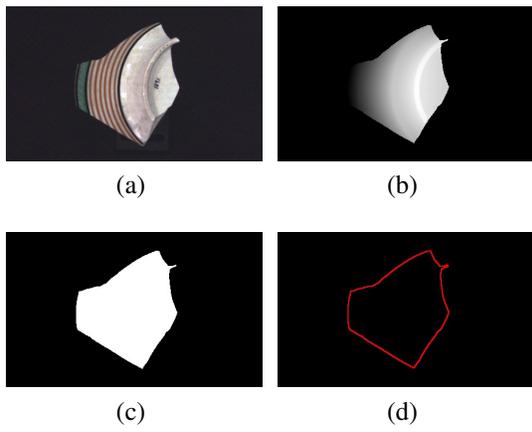


Figure 3

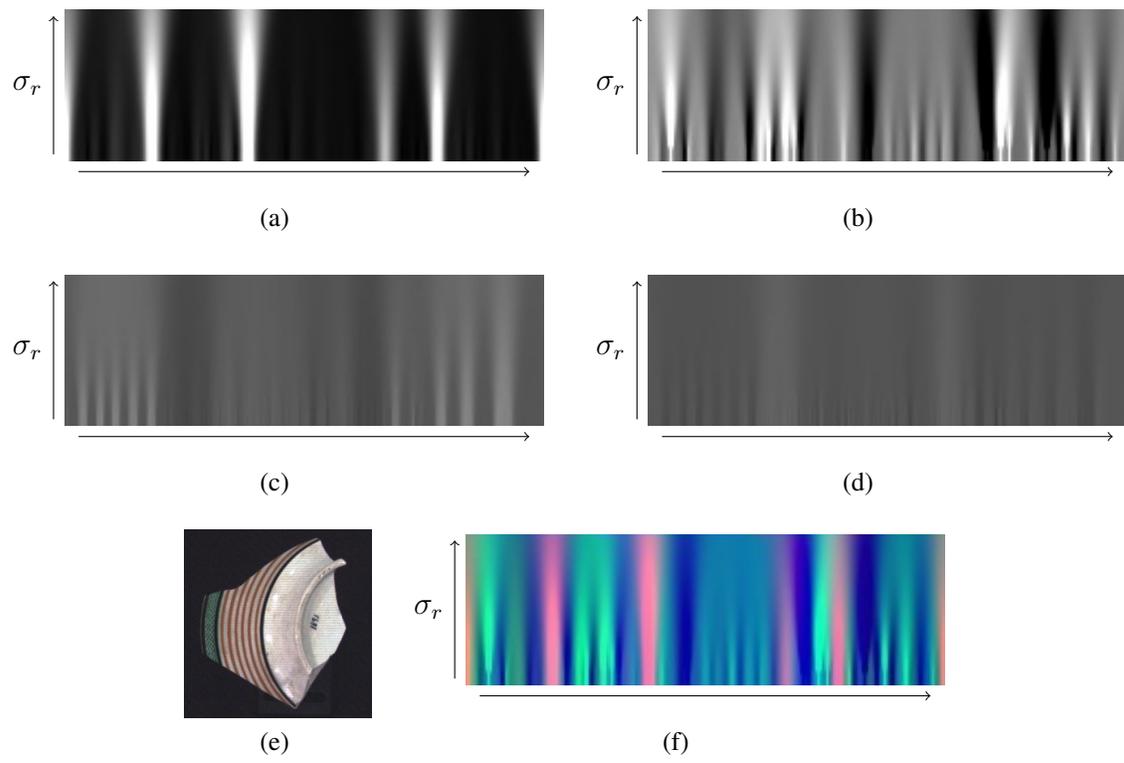
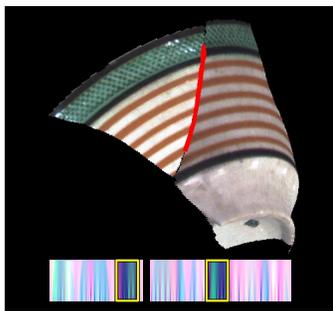
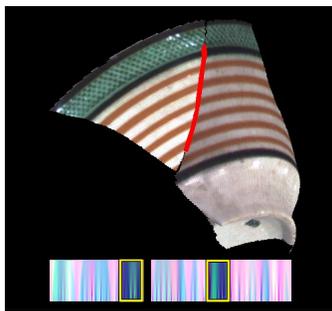


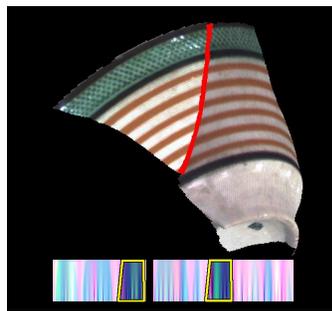
Figure 4



(a)



(b)



(c)

Figure 5

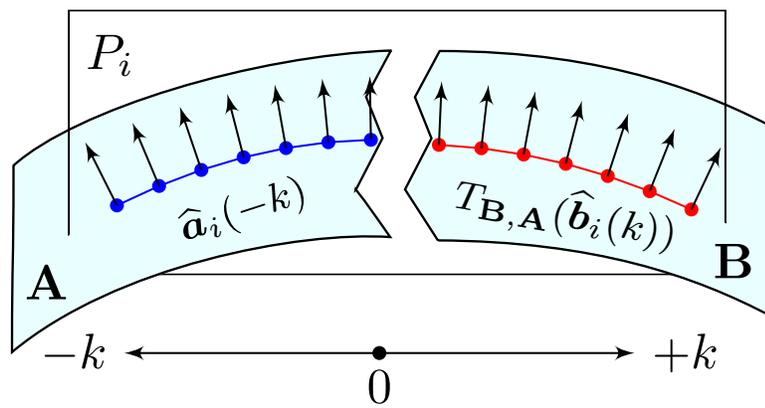
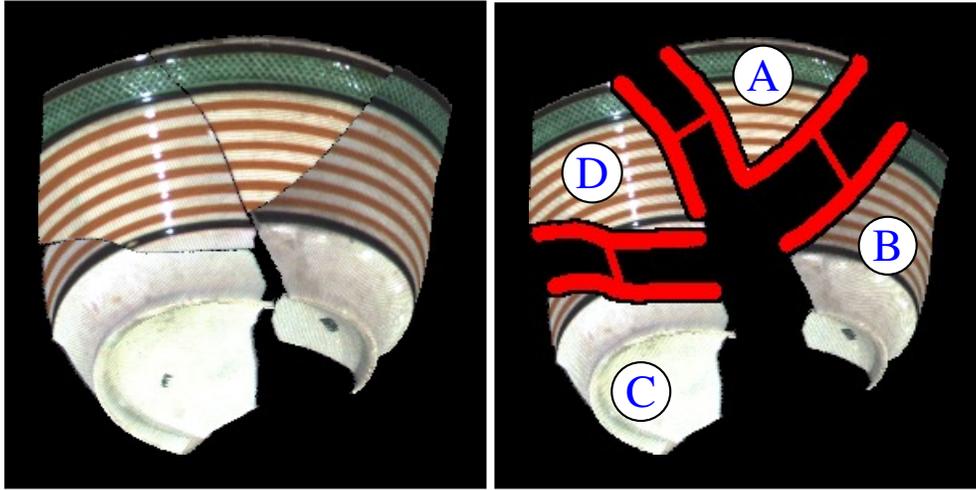
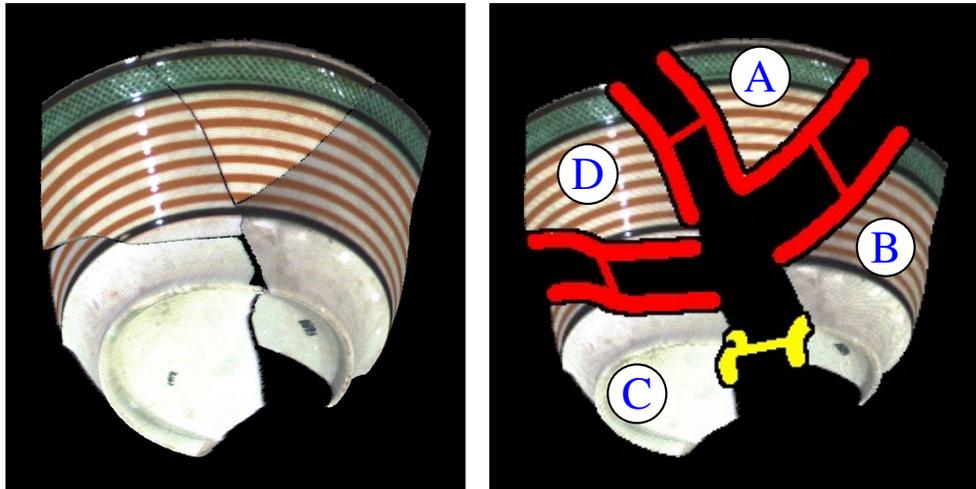


Figure 6

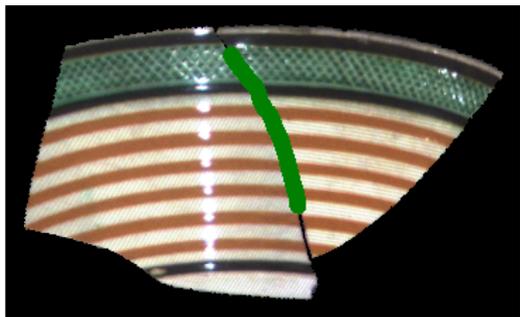


(a)

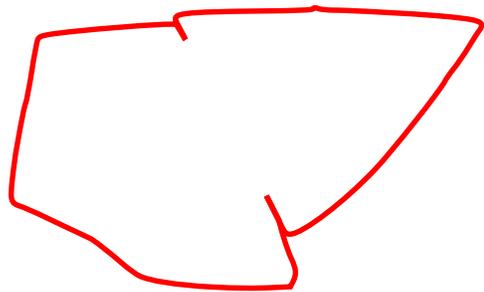


(b)

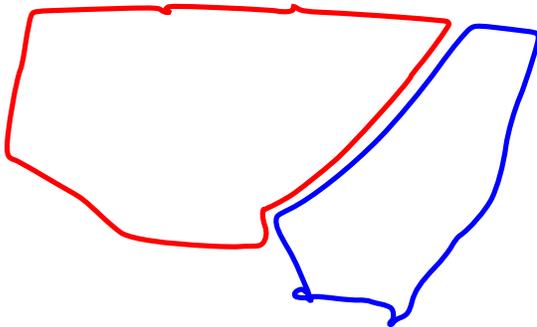
Figure 7



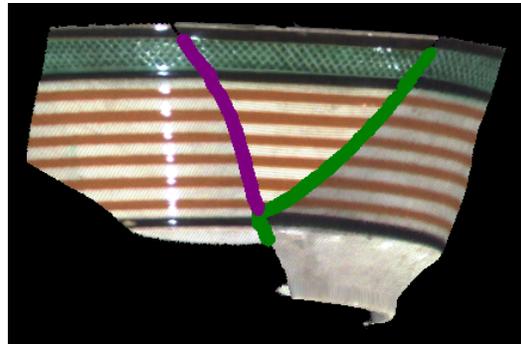
(a)



(b)



(c)



(d)

Figure 8



Figure 9

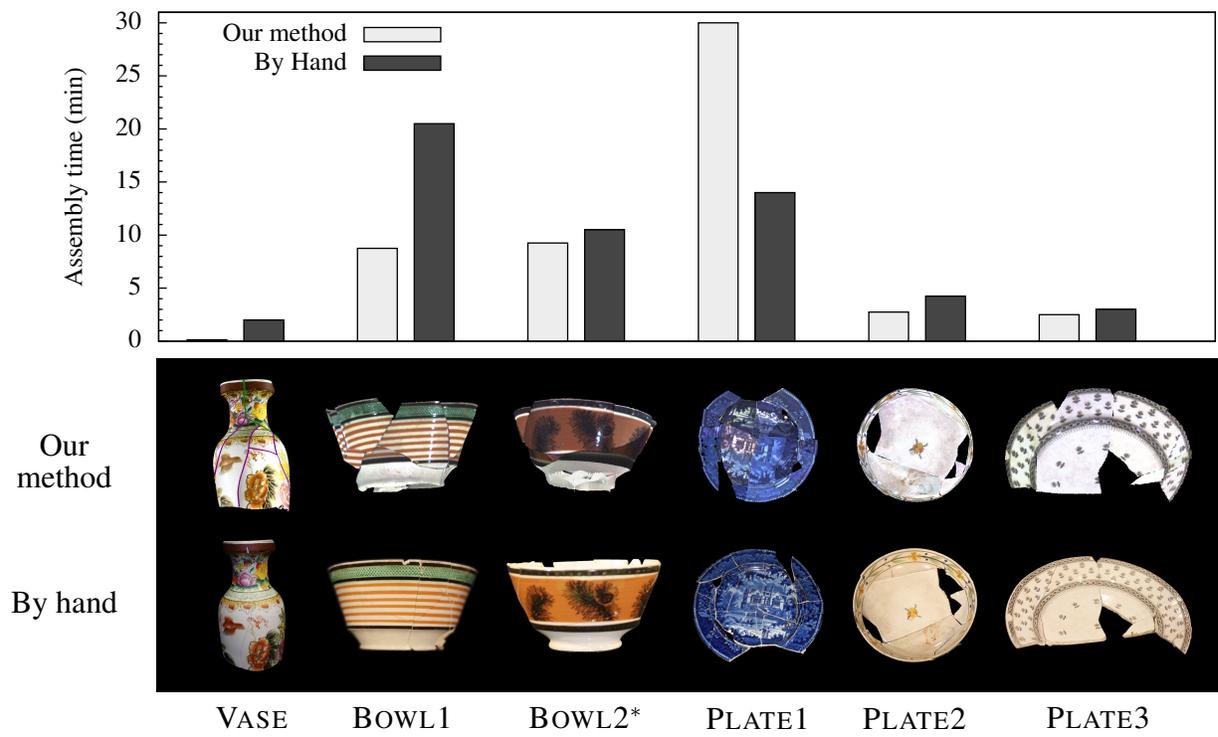


Figure 10

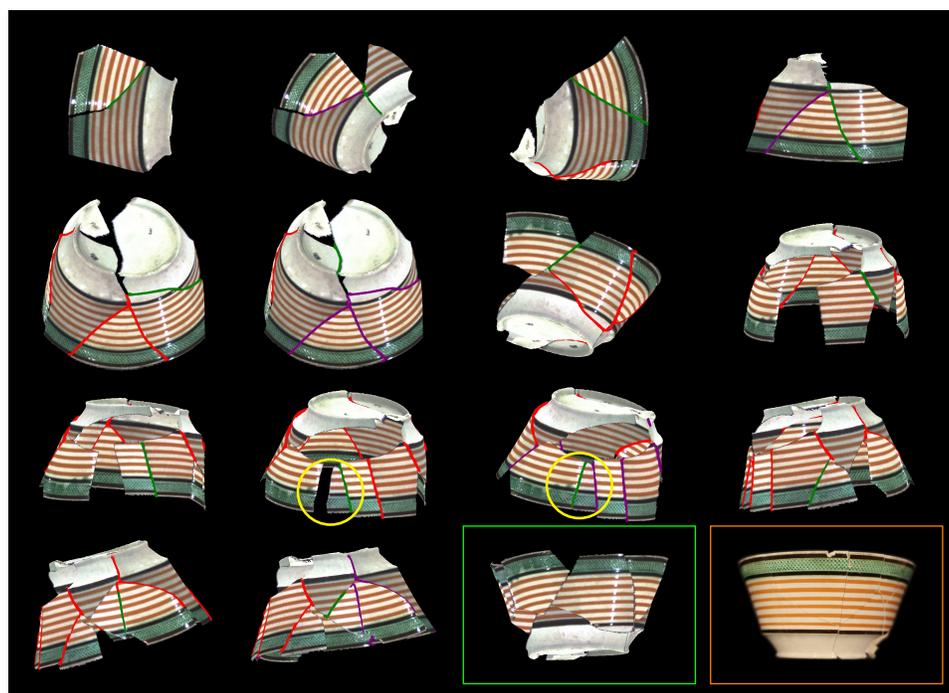


Figure 11

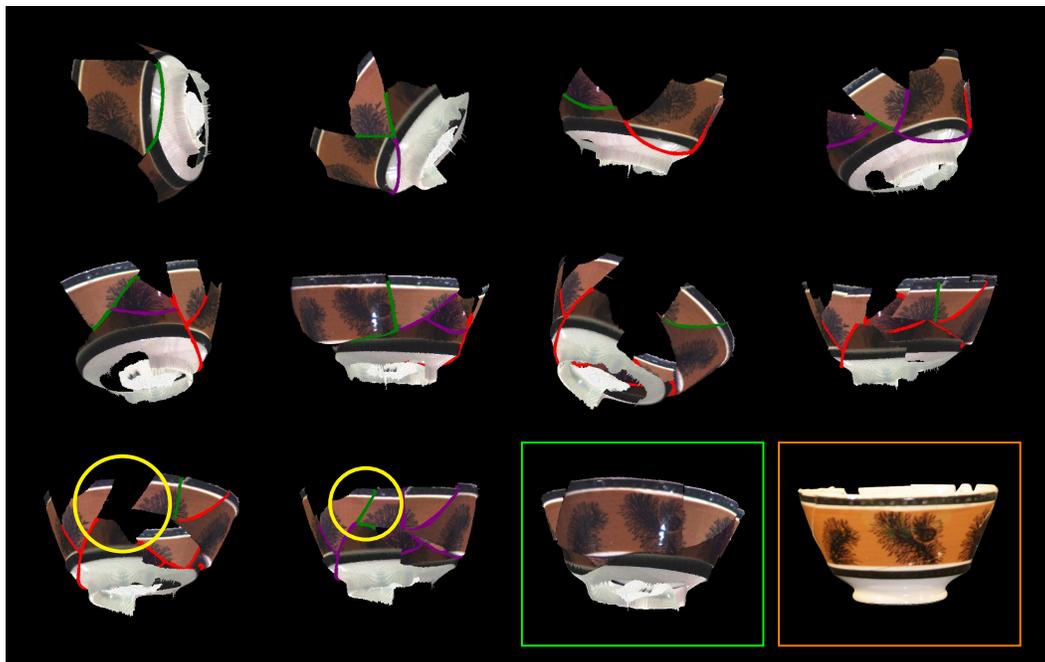


Figure 12

