Andrew R. Willis and David B. Cooper

Computational Reconstruction of Ancient Artifacts

From ruins to relics

Since the dawn of mankind, our ancestors have been creating man-made objects to serve in everyday life. Inevitably, these objects will erode, fragment, and break apart due to numerous causes which include so-called "acts of God" (such as floods and earthquakes), war, neglect, carelessness, and many more. During field excavations, archaeologists and anthropologists rarely find intact artifacts. More often, they find fragments of complete relics, and much work is spent studying individual fragments and trying to find correlations between the fragments. In these pursuits, they typically examine the geometry of the fragment and match the geometry of a given fragment against a candidate set of fragments that are hypothesized to have come from a single object. For flat surfaces such as stone tablets, this problem is similar to that of solving a large jigsaw puzzle where each piece is irregularly shaped. For free-form surfaces such as sculpture and pottery, the problem becomes a three-dimensional (3)-D jigsaw puzzle. More generally, use can be made not just of fragment geometry but also of fragment surface appearance since fragments often have natural or human-imposed patterns on their surfaces. In this article, we discuss the development of automatic reconstruction systems capable of coping with the realities of real-world geometric puzzles that anthropologists and archaeologists face on a daily basis. Such systems must do more than find matching fragments and subsequently align these matched fragments; these systems must be capable of simultaneously solving an unknown number of multiple puzzles where all of the puzzle pieces are mixed together in an unorganized pile and each puzzle may be missing an unknown number of its pieces.

SIGNIFICANCE OF ARTIFACT RECONSTRUCTION

Systems capable of automatically reconstructing objects from their fragments can greatly aid in the study of many civilizations and has potential applications in other domains such as forensic science, where they may be used to reconstruct broken bones or shattered glass and then estimate the energy incident to the glass structure and the point of impact. The key contributions of such programs are that they are huge time savers for archaeologists who may spend many hours assembling broken artifacts by hand. They also endow archaeologists with tools for quickly making precise measurements on fragments and also on reconstructed objects, thus permitting interpretations not heretofore possible. In fact, in most cases, broken vessels and artifacts

Digital Object Identifier 10.1109/MSP.2008.923101

1053-5888/08/$25.00©2008IEEE

are not reassembled unless the fragments are visually similar and discovered in a context that is local in both time and space. Automated reconstruction systems working from large databases of digitized fragments could uncover numerous partial or complete reconstructions of artifacts that may have been excavated during different years of the same excavation, or possibly from different sites altogether. In this way, reconstruction systems not only save researchers time but, given a sufficient database of fragments, also have the capacity to reconstruct artifacts that would have otherwise remained as an incoherent pile of disjoint fragments.

Ceramic pot fragments (referred to as sherds) are often the most important source of information from which archaeologists make inferences about the society that existed at a given site (see Figure 1). The primary use archaeologists make of the hundreds of thousands of sherds found at a site is to establish typological and chronological sequences of vessel types that developed during the course of a site's occupation. Indicative sherds, such as neck or base fragments or decorated pieces, allow them to infer the shape and type of the vessel in question. Indicative and nonindicative sherds are also counted to estimate the total number of vessel types used during a certain period in a specific location. Pottery sherds and reconstructed vessels are the best indicators of a site's chronology, the population's socioeconomic standards, and the extent of local or international trade. Estimating the vessel type, exact shape, size, and volume from a portion of the neck relies heavily on the completion offered by the classified vessel type. In such cases, subtle variations in estimated curvatures, especially for large vessels, may precipitate volume estimation errors of up to 30%. Such errors may be significant enough, if repetitive for one site or every site, to promote false typological and chronological conclusions. Completely or partially reconstructed vessels provide more accurate information on the exact vessels used on the site and can help prevent such difficulties. For example, Apollonia-Arsuf, a Mediterranean maritime city in Israel that flourished between the 6th century Before Common Era (BCE) and the 13th century Common Era (CE) was heavily involved in trade. One important aspect of this commerce involved the transportation of liquids in amphorae. The standard volume of these changed by about 10% from the 5th–6th century AD. Such subtle variations in volume are impossible to discriminate from only one or a few necks. Hence, complete or partial reconstructions provide a far more accurate assessment of the period that a specific type of amphora was in use and the numbers of these used. This in turn allows archaeologists to determine changing economic trends such as inflation or economic crises.

For tiles and murals that have figurative or even geometric and floral representations, more accurate reconstructions provide important information to archaeologists about the iconography and other explicit and implicit messages carried by the imagery, stylistic developments, etc. In the case of stone tablets and any kind of artifacts that have text on them, reconstructions that allow epigraphists to read the entire table text, sentences, phrases, or even individiual words provide new insights on the site from the stance and style of the letters, inferred dates, and other important information.

Among projects of significance on object reconstruction, two are the Digital Forma Urbis Romae project [2] and the Stitch project [3]. The Digital Forma Urbis Romae project seeks to

(a)

(b)

(c)

[FIG1] Here is an example of the large number of ceramic fragments found at a typical archaeological site. Near the city Herzliya-Pituach, Israel, the Apollonia-Arsuf site contains a multitude of Greek, Byzantine, Muslim, and Crusader ceramics. Some finds such as those in (a) are too numerous to store in museums and must remain on site, while others such as those in (b) and (c) have more archaeological significance and are painstakingly assembled by hand off-site.

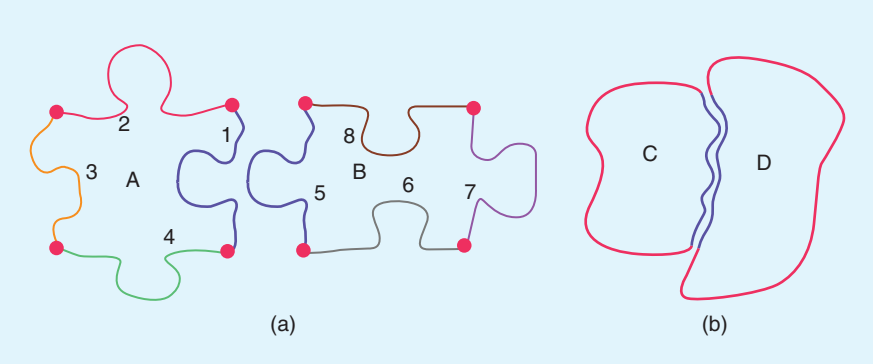

[FIG2] The differences between reassembling commercial jigsaw puzzles and reconstructing broken artifacts. (a) Jigsaw pieces have readily identifiable corners (red dots) allowing programs to easily separate portions of the boundary (shown as different color curves) that will match with some other unique puzzle piece. Additionally, each boundary segment is a smooth planar curve having an isthmus, or neck, that is highly indicative for finding that unique matching puzzle piece. (b) Two hypothetical fragment boundaries. Note here that the problem is made much more difficult as corners are not easily identifiable and may not indicate the beginning or end of a curve that will uniquely match some other fragment. Worse, any portion of a boundary curve may match to any other fragment and the curve itself may match equally well with numerous similar boundaries from other fragments. (Used with permission from [1].)

reconstruct a giant marble map of ancient Rome dating to the reign of Septimius Severus (circa 200 AD) where only a small portion of fragments from the original map remain (1,186 of them remain), which make up a fraction of the original map. The Stitch project seeks to develop a generic system capable of recovering the shape of broken pottery vessels from measurements of their fragments. As mentioned previously, ceramic pottery fragments are commonly the most prolific find at an archaeological site, and they are a major source of information concerning the functioning of the society at the site. A system capable of automatically sorting through and matching pottery fragments is not only a significant time saver but also provides new information tools for archaeologists and anthropologists that may be leveraged to improve our understanding of ancient civilizations.

Are these systems doing signal processing? Yes, these problems are the most challenging form of signal restoration from chaotic noisy data. Here the signals are combinations of multiple planar or 3-D curves and 3-D surfaces. The chaos is manifested as the signals appearing in collections of unorganized fragments and the noise as deformations of the signals plus random perturbations arising in the measurement process. Perhaps we think of this as decomposing a complex data set (image) into patches that we then deform with distortions and noise, and then forget the side information that tells us how to reassemble the patches. Generally, we do not just have fragment shape geometry to work with. There may also be patterns on the surfaces of the fragments.

Initial work on puzzle solving may be traced back to the problem of computational reconstruction of jigsaw puzzles, which has been discussed as far back as the mid-1960s [4] with some additional work in the 1980s [5], [6]. However, numerous works on the problem have emerged in recent years, and current systems are capable of automatically assembling jigsaw puzzles with hundreds of pieces [7], [8].

The problem of reassembling real-world archaeological puzzle pieces, i.e., reassembling free-form broken objects from measurements of their fragments, has received little attention until recently. As shown in Figure 2, this problem is more difficult than assembling commercially produced jigsaw puzzles. A major complication is that 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 objects; 3) fragments may be missing as they may have yet to be discovered or may have been destroyed by some physical phenomenon.

PROBLEM STATEMENT AND TERMINOLOGY

We frame the problem of reconstructing artifacts as a generic geometric estimation problem that is formalized when we discuss specific approaches to this problem proposed in the literature. Initial discussions emphasize only the problem of estimating the geometry of a single object from its fragments as this controls the scope of the problem considerably and is the emphasis of current approaches. We then broaden our scope of consideration to include approaches that use other characteristics such as patterns on the surface to either influence the geometric reconstruction result [1] or as the basis of reconstruction [13], [14].

We begin by denoting the global geometry of the fractured object as a large collection of unknown variables, Ω, that may be decomposed into two subsets: 1) Ω_{outer}, those variables that characterize the *outer boundary* of the object before being fractured, i.e., the outer boundary, and 2) Ω_{inner}, those variables that characterize the boundaries generated when the object was fractured, i.e., the *fracture boundaries*. A third set of variables describes locations where three or more fragment boundaries have broken apart and are referred to as *fracture junctions* or simply *junctions* and denoted $\Omega_{\text{inner}\perp}$ (see Figure 3). A fourth set of variables describes special junctions between the intact outer boundary and the fracture boundary and are referred to as *outer-surface junctions* and denoted $\Omega_{\text{outer}\perp}$. Together these variables account for all the usable geometric information for geometric reconstruction techniques. Measurements from fragments allow us to estimate some subset of these unknown vari-

ables. Our task in geometric reconstruction is to define a computational method to merge these estimates to uncover not only the shape of the global intact outer surface but also how that surface was fractured when it broke apart in antiquity by computing an estimate for Ω.

For clarity, we have divided existing approaches into two classes based on the underlying geometry which each method seeks to estimate: i) two dimensional (2-D)-reconstruction, i.e., the reconstruction of flat objects such as stone tablets or frescoes, and ii) 3-D-reconstruction, i.e., the reconstruction of 3-D objects. The primary difference between these approaches is the form of the geometric primitives that must be matched. In fact, Table 1 specifies the differences in these primitives and specifies how the generic terminology used for the problem statement changes when we discuss 2-D and 3-D reconstruction. When reviewing this table, note that typically Ω_{outer} and Ω_{inner} are a single contiguous planar curve or surface patch. However, in some cases, a single fragment may have two distinct outer surface patches such as that in the leftmost image of Figure 4. Here, the cyan-colored surface regions show two disjoint surface patches that together define Ω_{outer}.

COMMON SIMPLIFICATIONS

CONVERT THE PROBLEM TO 2-D

All real-world objects considered for reconstruction are 3-D. Yet, for flat objects such as stone tablets, a simplification is made to convert the problem to that of matching 2-D planar curves by using the tablet fragment outlines as viewed from either the top or bottom. As a convention, we assume the top surface is used for reconstruction. For thin tablets or tablets that have uninformative break surface geometry, this simplification can reduce computational complexity without adversely impacting the reconstruction accuracy [10]. Theoretically, only the tablet thickness information is lost if the tablet has a slab geometry and is extremely thin. In this case, the entire 3-D geometry of the tablet may be generated by extruding both the intact and fracture boundaries along the tablet plane normal by an amount equal to the thickness of the tablet, i.e., the entire 3-D surface is an explicit function of the thickness and the estimated 2-D planar variables.

Use fracture junctions—Several researchers [1], [15]–[17] have noted that special locations where three fragments meet are useful for reconstruction. These locations are referred to as "T" or "Y" junctions depending on the relative angles between the three incoming fracture boundaries. In [1], the authors report that for their experimental tablets, between 70% and 90% of all interior junctions were "T" junctions, and between 6% and 10% of the remaining junctions are "Y" junctions. The utility of detecting these locations for reconstruction is shown in Figure 3, and some methods seek to identify these locations for boundary matching. While higher-order junctions (junctions involving >3 fragments) occur, their relative frequency of occurrence is small from empirical observations.

Dispose of the intact boundaries—Most reconstruction approaches also either classify the outer boundary automatically or semi-automatically by assuming the surface texture is different on the intact and fracture surfaces as in [9] and [11], or by assuming that the boundary has been manually identified or otherwise removed [1], [10] [see Figure 3(a) and (b)].

METHODOLOGY

Computational methods for object reconstruction must have

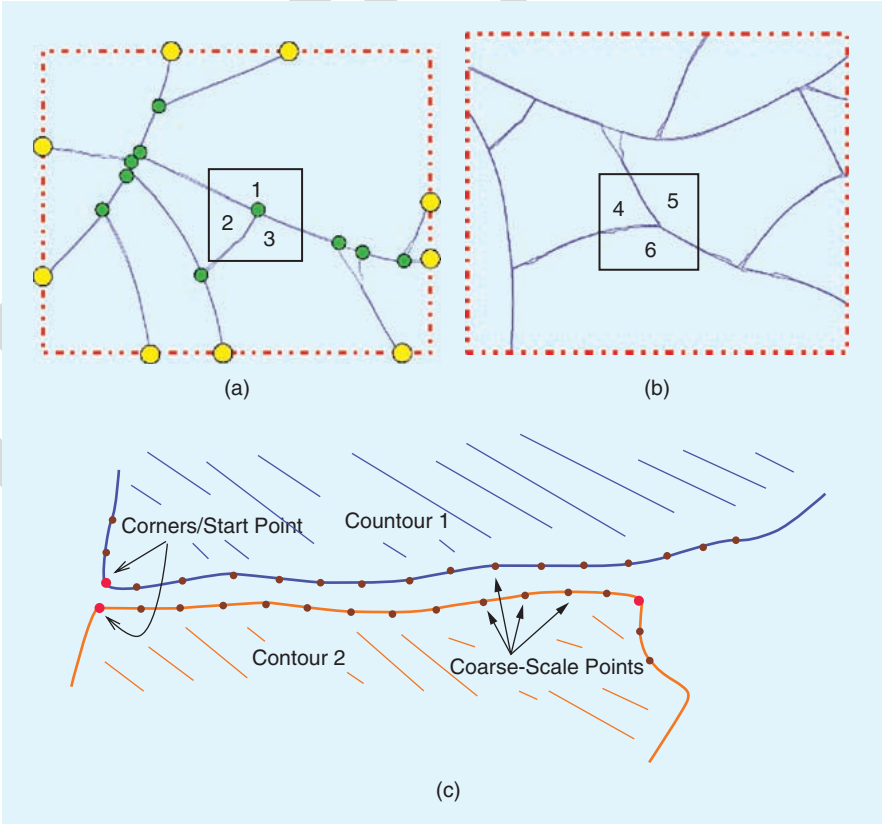

[FIG3] In these modified figures from [1], we define the geometric features used by various approaches for assembly. In (a), the outer contour of the square ceramic tile, i.e., the intact boundary, Ω_{outer}, is shown in red; fracture boundaries, Ω_{inner}, are shown in blue; vertices, $\Omega_{\text{inner}\perp}$, are shown as green points; and outer vertices, $\Omega_{\text{outer}\perp}$, are shown as yellow points. Note that most automatic assembly algorithms including [2], and [9]–[12], only consider the fracture boundaries, Ω_{inner}, and perhaps the vertices, $\Omega_{\text{inner}\perp}$. The black boxes in (a) and (b) show "T" and "Y" junctions, respectively, that ceramics tend to generate when fractured. These junctions are typically high curvature locations along the fragment fracture boundary and are important for (c) reconstruction as they often denote origin points for fracture boundaries that match to only one other fragment.

digital representations for the fragments to be assembled. There are a large variety of potential techniques for obtaining 2-D and 3-D fragment measurements. Measurement technologies for ancient artifacts are almost all noncontact, i.e., optical, measurement systems. For 2-D reconstruction, digitization is commonly done using rectified images of the tablet fragments, i.e., images that provide metric data for fragment boundaries (see Figures 16 and 12 for examples). For 3-D reconstruction, digitization is usually accomplished by commercially available 3-D laser-scanners that now typically produce highly accurate measurements with (x, y, z) errors <0.25 mm (see Figures 4, 6, 10, 11, 14, and 15 for examples). Yet, there are numerous image-based techniques from computer vision literature that capture 2-D and 3-D geometry from images. Popular methods here include stereoscopic reconstruction [18], photometric stereo (shape-from-shading) [19], and 3-D surfaces from occluding contours in many calibrated images [52], [53]. The ultimate output of these measurement devices is a digital representation, i.e., digitized version, of the fragment as either a 2-D boundary curve or 3-D surface mesh.

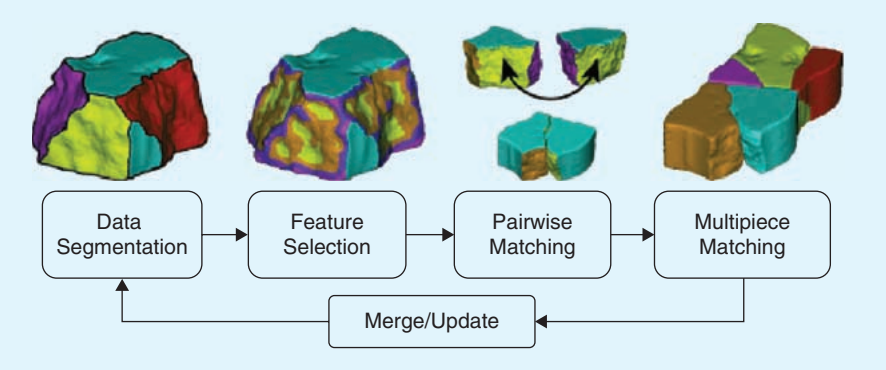

[FIG4] The generic process for object reconstruction. It consists of four steps (from left to right): 1) classify the boundaries of interest for matching, 2) compute a shape model for these boundaries in terms of features, 3) search for correct pairwise matches of the computed features and 4) implement a merge and update procedure that merges pairwise matches into larger multipiece configurations. (Used with permission from [11].)

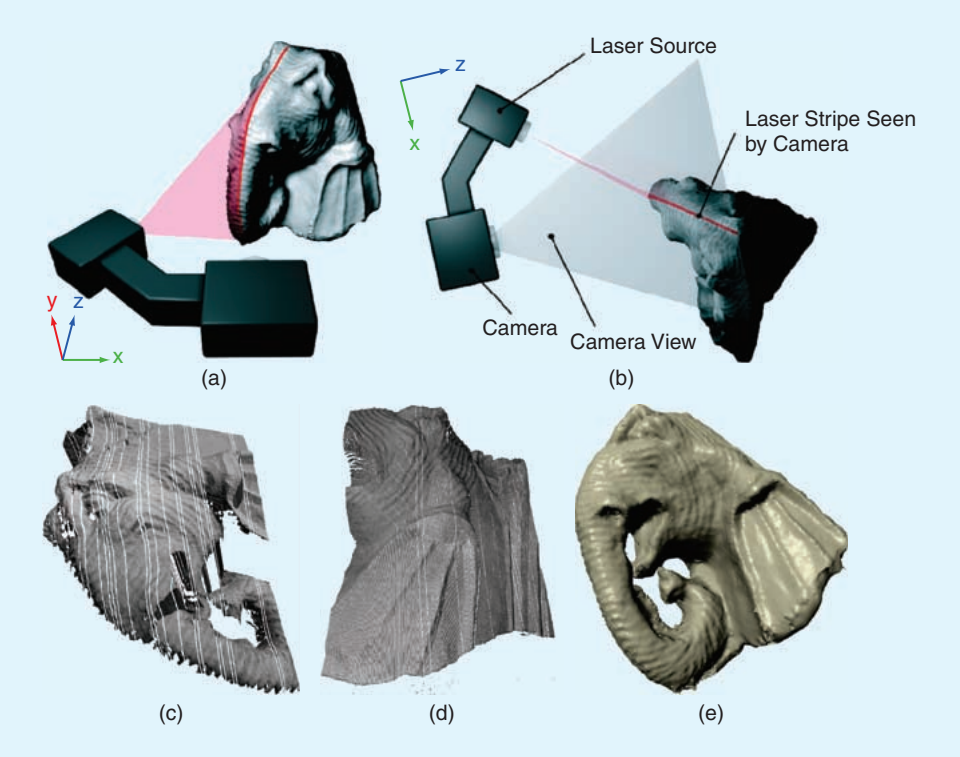

[FIG5] (a), (b) Views of a typical laser scanner used for artifact digitization. A stripe of laser light is swept across the surface of the object, and the laser stripe projected onto the object surface is observed by an imaging sensor. The sensed locations of the observed laser stripe allow for quick computation of a range image, $z = f(x, y)$. By changing the relative position of the sensor and object, multiple range images may be measured, as shown in (c) and (d), that are subsequently merged to generate a (e) complete model of the 3-D object .

While measurements for a complete 2-D fragment can usually be captured with one or two images from a calibrated camera, generating a complete model of a 3-D fragment is often laborious and time consuming. Laser scanning technologies typically recover surface depth measurements through triangulation using a laser line emitter and high-resolution image sensor with known relative positions and orientations [see Figure 5(a) and (b)] [20]. The result is a range image, i.e., a collection of measurements that have the form $z = f(x, y)$, where z is the object depth for an (x, y) location in the sensor image plane. By changing the relative orientation of the sensor and fragment, different range images are captured, each of which provides a partial description of the fragment surface [Figure 5(c) and (d)]. The complete fragment model is generated by aligning the measured range images to a common coordinate system using techniques such as [21]–[23] and subsequently merging the range images using techniques such as [24] to produce a virtual 3-D model of the real-

world fragment [Figure 5(e)].

Using digitized fragments as raw data, assembly approaches follow a four-step process as summarized in Figure 4 to reassemble the fragments. Raw data for a typical 2-D fragment boundary consists of 1k–5k (x, y) points and a typical 3-D fragment surface may have 0.5M–2M (x, y, z) points. Considering a single broken artifact may consist of 10–30 fragments, direct manipulation of the raw data has a huge computational cost. Reconstruction approaches typically start by subdividing the fragment into matchable subsections, i.e., segmenting the boundary into regions to be matched (Figure 4, step 1). For each identified boundary subsection, a compact representation of the fragment shape is computed in terms of features taken from the fragment surface (Figure 4, step 2). Estimated features from portions of each fragment are then compared against other fragments to identify candidate pairwise matches between fragments (Figure 4, step 3). Finally, an assembly rationale is enforced by a search procedure which updates and merges pairwise matches into large configurations of aligned fragments until a stopping criterion is met (Figure 4, step 4).

As the fragment matching is geometric, features computed in step 2 often originate from classical literature on differential geometry such as [25], and there is a considerable body of work within the computer-aided design and computer graphics community that discusses applications of these classical theories to discrete curves and polygonal surfaces (see Figure 6). Many approaches use intrinsic shape parameters that provide Euclidean-invariant representations for the fragment geometry [10]–[12], [15], [26]. A huge benefit of using such invariant features is that it allows researchers to define performance functionals that directly compare locally computed features on each fragment to determine fragment-to-fragment correspondences [27].

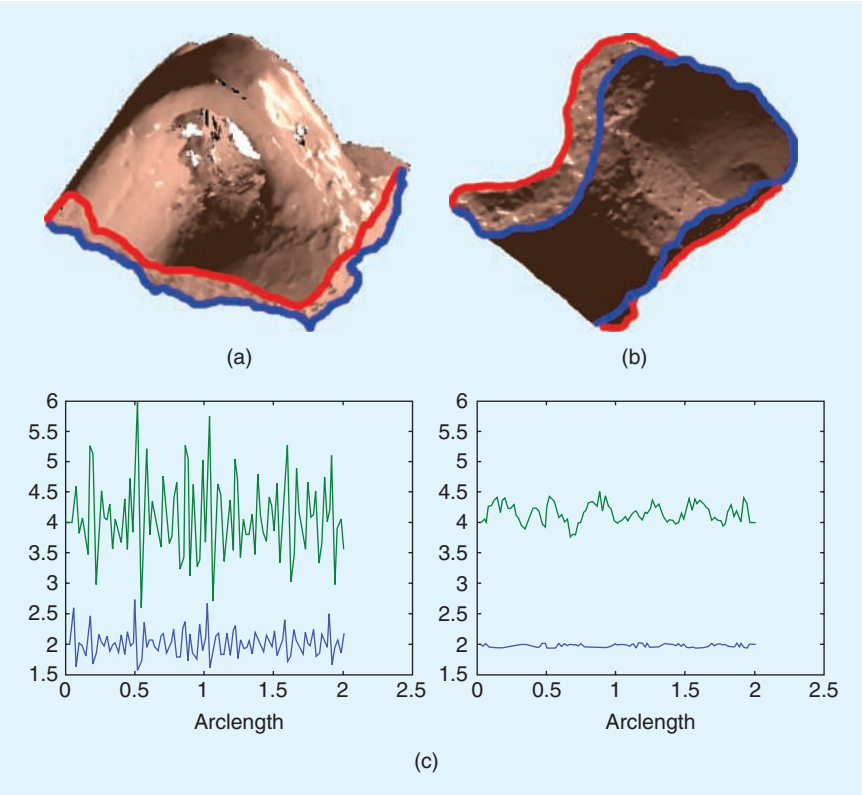

[FIG6] In these figures from [12] and [28], Differential properties of fragment boundary curves such as curvature and torsion are used to reconstruct both 2-D tablets (curvature-only models) and 3-D ceramic pottery. (a) and (b) views of two pot fragments and that portion of the outer break curves visible from these views in red and blue. In [12], [28] a subset of this complete outer break curve, $\Omega_{\text{outer}\perp}$, is used as their parameters for reconstruction by elastically matching the curvature and torsion signature of the blue curves. (c) The curvature signature (in blue) and torsion signature (in green) for a 3-D space curve using two different estimation methods for these differential properties. Finite difference methods are shown on the left and essentially nonoscillatory (ENO) methods are shown on the right. Reference [28] reports that the ENO method controls noise in the estimate and provides improved accuracy for 3-D space-curve matching. (Used with permission from [12] and [28].)

APPROACHES FOR 2-D RECONSTRUCTION

Two approaches representative of the state-of-the-art in computational tablet reconstruction are [10] and [1]. These approaches follow the generic assembly process outlined in Figure 4 with the exception that the authors of [10] do not subdivide the fracture curve into segments using vertices as indicated in Figure 3. Subsequent differences between these approaches are primarily: i) the features selected for matching, ii) the curve-matching performance functional, and iii) how correct matches are detected and merged into larger puzzle solutions. While not focused on archaeological reconstruction, we also wish to mention work by forensic researchers [29] that propose a related approach for reconstructing shredded paper documents.

FEATURE SELECTION

In [10], only fragment fracture curves, Ω_{inner}, are used for reconstruction. These curves are represented in terms of curvature as a function of the curve arc length. For discretely sampled boundaries, this representation is obtained by computing local curvatures of the fragment boundary when sampled at uniform intervals of arc-length. We refer to the resulting discrete one-dimensional function as the curvature signature, $\kappa(s)$, which has been extensively studied for curve matching with seminal work on the topic available in [30].

Provided that samples are taken sufficiently dense, this representation benefits from two theorems of differential geometry: i) the fundamental existence and uniqueness theorem, which ensures that, for planar curves, $\kappa(s)$ is completely equivalent to the original boundary curve up to an unknown position and ori-

entation [31], and ii) two planar curves are congruent if and only if they have the same curvature signature $\kappa(s)$ [25]. Since the curvature signature is an intrinsic property of the boundary, it is coordinate free, i.e., it is invariant to Euclidean transformation. This allows matching fragment boundary shapes by directly matching their curvature signatures.

Unfortunately, the position- and orientation-free representation afforded by the curvature signature is countered by the inherently noisy values obtained when computing curvatures, especially for fragment boundaries. Problems in obtaining accurate curvature values can be attributed to 1) the curvature depends upon the second-order derivative of the fragment boundary and 2) fragment boundaries typically delineate regions where the surface derivative is locally discontinuous; measurements in these areas are more variable and less accurate. These effects are widely acknowledged in the literature, and nearly all techniques using the curvature signature (or other differential quantities such as torsion) perform low-pass filtering, i.e., curve smoothing, to control the estimated curvature variability [10], [12], [27], [28], [30], [32]. Another important source of variability is any bias present in the estimation technique. Details on estimation approaches and their relative merits are widely available in both 2-D [33], [34] and 3-D [34]–[39].

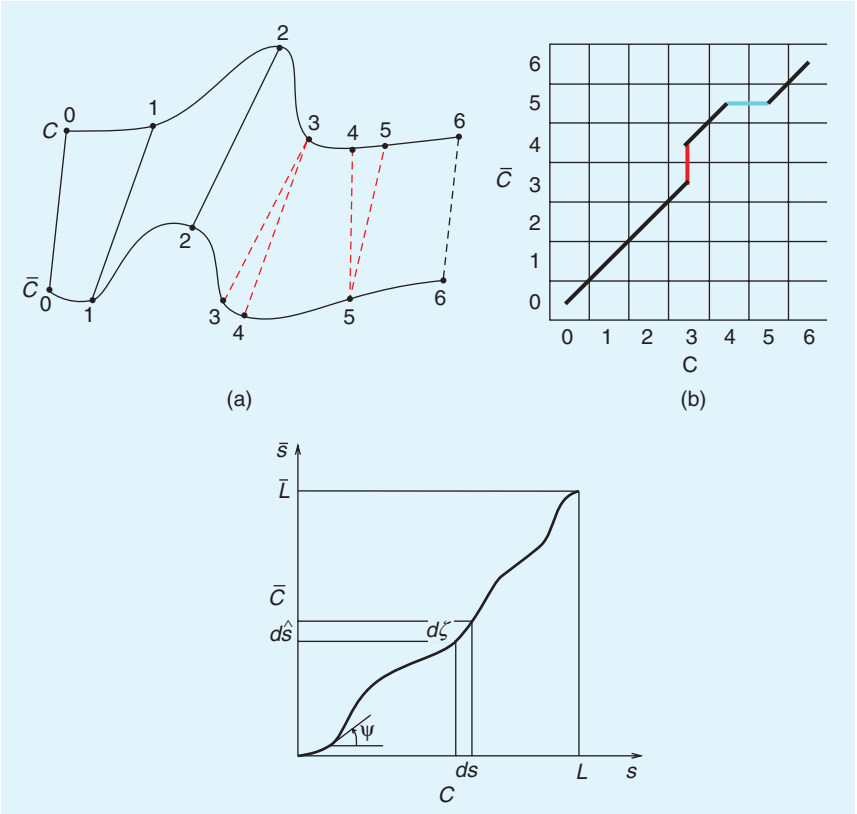

[FIG7] An example of elastic curve matching between the curves C and $\bar{C}$ starting at the initial index 0. Due to differences in the curve shape and sampling, the elastic matching suggests a correspondence that is not one-to-one. This correspondence is specified as a sequence of corresponding sample indices (i, j) where the ith sample in C is said to correspond to the jth sample in $\bar{C}$. For the curve in (a), the $(C, \bar{C})$ correspondence pairs are $(0, 0), (1, 1), (2, 2), (3, 3), (3, 4), (4, 5), (5, 5), (6, 6)$. Dynamic programming is often used to find the optimal, i.e., lowest cost, correspondence between the two curves that is often visualized as a curve in the correspondence space, referred to as a *correspondence curve*. (b) The correspondence curve for the curve match shown in (a). Black regions of the path indicate one-to-one correspondences, the red region indicates where the curve $\bar{C}$ has compressed to match a smaller portion of arc length on C and the blue region indicates where the curve $\bar{C}$ has stretched to match a larger portion of arc length on C. (c) A continuous correspondence curve, ξ, used to define the edit distance matching metric used in [1] [(see (1) and associated discussion for details)].(Used with permission from [28].)

In contrast to [1], [10] uses fragment fracture curves, Ω_{inner}, and vertices, $\Omega_{\text{inner}\perp}$, to perform reconstruction. As in [10], these curves are represented by curvature signatures for the large majority of curve comparisons. Yet, their method also relies on the measured 2-D (x, y) boundary points in the final matching stages to determine the best matching candidate from a small number of highly similar candidates.

Given these features as a representation for 2-D boundary curves, the two approaches must identify matching fracture curve segments from corresponding fragments. Each curve segment is a contiguous subset of the complete fragment fracture curve and has an unknown length that is the arc length of the curve generated when the pair of corresponding fragments broke apart. In [10], both the correspondence and length of these fracture curve segments are estimated as part of the curve matching process that allows curves to be matched along any portion of the complete fracture curve. In contrast, [1] searches through the smaller set of detected fragment vertices, $\Omega_{\text{inner}\perp}$, to identify matching fracture curve segments (see Figure 3), which involves significantly less computation.

MATCHING PERFORMANCE FUNCTIONAL

As archaeological fragments are often weathered or chipped, both the arc-length and the measured differential quantities of the curves are very noisy and will not exactly agree. Hence, classical curve matching techniques such as cross correlation and least squares are prone to failure. For this reason, archaeological curve matching approaches use elastic curve matching methods as described in Figure 7, which have shown to be much more robust. Such methods have a long history in a wide variety of other applications such as biological-sequence matching, dense pixel matching between a pair of rectified stereo images, and speech phoneme recognition.

As shown in Figure 7, elastic-curve matching defines a correspondence between two curves as a collection of sample index pairs, e.g., the integer pair (i, j) indicates that sample index i

from fragment boundary C corresponds to sample index j on fragment boundary $\overline{C}$. A sequence of pairs then specifies a point-by-point correspondence between a pair of fragment boundaries and may allow for a two-or-more to one mapping between curve points, which is necessary when dealing with archaeological fragments.

Reference [10] defines a a curve matching performance functional that seeks to enforce two soft constraints: i) corresponding points on the boundary should have the same curvatures and ii) the correspondence between the curves should be as close to one-to-one as possible. The first constraint is enforced by a weighted distance metric consisting of the sum of squared differences of the matched curves at corresponding locations and weights this difference by an approximation of arc-length obtained by taking the sequential difference in the matched curve sample indices. The second constraint is enforced by measuring the Manhattan distance between the correspondence curve in Figure 7(b) and the correspondence curve defined by a one-to-one correspondence, i.e., a straight diagonal line.

As mentioned previously, [1] uses two curve matching performance functionals: i) curvature signature matching and ii) a Euclidean matching for fine detail. For i) the authors define a performance functional which they call edit distance. This distance metric is derived from a variational formulation that seeks to minimize the stretching and bending energies necessary to match two different curves. These energies are defined on a continuous correspondence curve ξ that maps the curve C to $\overline{C}$ using the variational equation (1) as shown in Figure 7(c)

$$\mu(\psi) = \int_0^L |\cos\psi - \sin\psi| + \lambda|\kappa(h)\cos\psi - \kappa(\bar{h})\sin\psi| d\xi, \qquad (1)$$

where ψ denotes the angle of the curve ξ with respect to the correspondence-space x-axis. Note that if $\psi = \pi/4$, the first term, the stretching energy, is zero which corresponds to a one-to-one matching of the curves C and $\overline{C}$. The second term, the bending energy, is a measure of how much the matched curves must bend in order for their curvature values to agree given the matching specified by the correspondence curve ξ. Since different regions (in terms of arc length) along each curve may correspond, separate parameters h and $\bar{h}$ are defined that denote matched locations along C and $\overline{C}$, respectively. A mixing parameter λ, is introduced to control the relative influence of the two terms in (1). They then may solve the curve-matching problem by finding the curve ξ, or equivalently, the sequence of angles ψ originating from the origin that minimizes the total integrated energy in (1). The authors have developed a dynamic programming approach to efficiently solve the discrete version of this problem to compute curve matches.

For ii), the authors use another dynamic programming curve matching technique to obtain correspondence between the measured (x, y) boundary points. Using this correspondence, techniques from [40] and [41] are used to quickly compute the Euclidean transformation that minimizes the squared distances between corresponding points, aligning the two fragment boundaries. For this purpose, the performance functional for curve matching is a sum of three terms weighted by hand-specified parameters: 1) the observed squared error between the matched points after alignment, 2) the length of the matched curve, and 3) the amount of variation in the boundary, i.e., a measure of the complexity or jaggedness of the matched curves.

COMPUTATIONAL CONSIDERATIONS

Both [10] and [1] use a coarse-to-fine (multiscale) approach to improve performance by quickly reducing the search space for compatible matches. Coarse representations of the fragment fracture curve may be compared at lower computational cost and serve to quickly identify matches that are clearly incorrect. Progressively higher detail representations are then used to refine the set of candidate matches for each fracture curve segment until a final set of candidates are identified. The multiscale approach is implemented by low-pass filtering the curvature signature of the fragment boundaries using a Gaussian filter bank whose cutoff frequency increases as the scale, or equivalently, the level of detail, increases. Lower frequency representations of the boundary require less samples and may be quickly compared to eliminate a large portion of the candidate fragments that do not match the fragment of consideration.

SEARCHING FOR THE SOLUTION

In [10], the end goal of the automated system is to identify some matches within the puzzle. Hence, there is no governing constraint that the system provide a global solution. In fact, it is troublesome to define the existence of such a solution given that, in a generic situation, the system may be missing fragments and, in some cases, include pieces from other unrelated objects. With this perspective in mind, the authors of [10] specify a multiscale procedure that uses their curve matching performance functional to reduce the number of candidate matches at each scale until a unique fragment match is obtained for some contiguous subset of each fragment boundary. At this point, the pairwise correspondences provided by matching the intrinsic properties of the fragments allows fragment boundaries to be aligned within a single global coordinate system that is the puzzle solution.

In [1], the authors use a variation on the best-first strategy which they call best global first. This technique looks at "T" and "Y" junctions and seeks to match fragment boundaries at open vertices, i.e., vertices where less than three pieces have been matched. The curve-matching performance functional is applied to find a short list of possible matching candidate fragments. Each candidate in this short list is assigned a global confidence score that favors fragments that fill space at open vertices. The fragment with the best combined curve-match and global confidence score is then added to the configuration of pieces. To account for potential incorrect matches, k different solutions are considered in parallel by starting the search process with a different initial fragment. The result is a set of k solutions to the

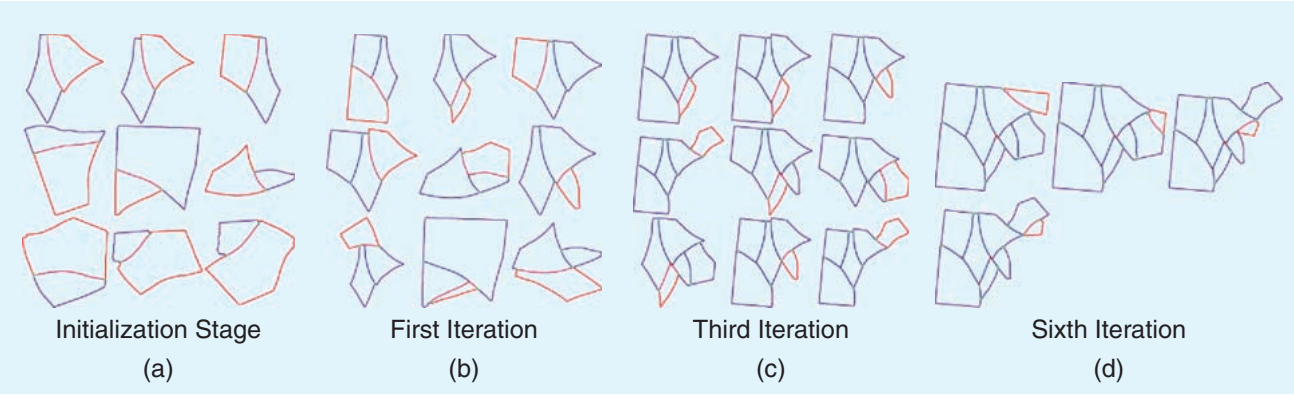

[FIG8] Nine fragment pairs, shown in (a), are chosen as different starting points for reconstructing a ceramic tile. The results of the best-global-first search for matching fragments are shown at iterations (b) 1, (c) 3, and (d) 6. (Used with permission from [1].)

puzzle which may or may not agree on the final global solution (see Figure 8).

APPROACHES FOR 3-D RECONSTRUCTION

While more than ten different approaches exist for reconstructing 3-D objects, we restrict our attention to four approaches that represent current state-of-the-art for 3-D object reassembly [9], [11], [23], and [16], which is a recent variant on the approach [43]. Only two of the four discussed approaches, [11] and [9], are designed for reconstructing generic free-form 3-D geometry. The approach in [43] is focused on reconstruction of axially symmetric surfaces, i.e., ceramic pots made on a potter's wheel, and [23] is not a complete reassembly system but rather a collection of experimental alignment techniques applied for aligning incised tablets from a marble map of Rome: the Severan Marble Plan of Rome, or the Forma Urbis Roma. With these major differences in mind, we begin this section by discussing approaches [43] separately and then we discuss [11] and [9] together.

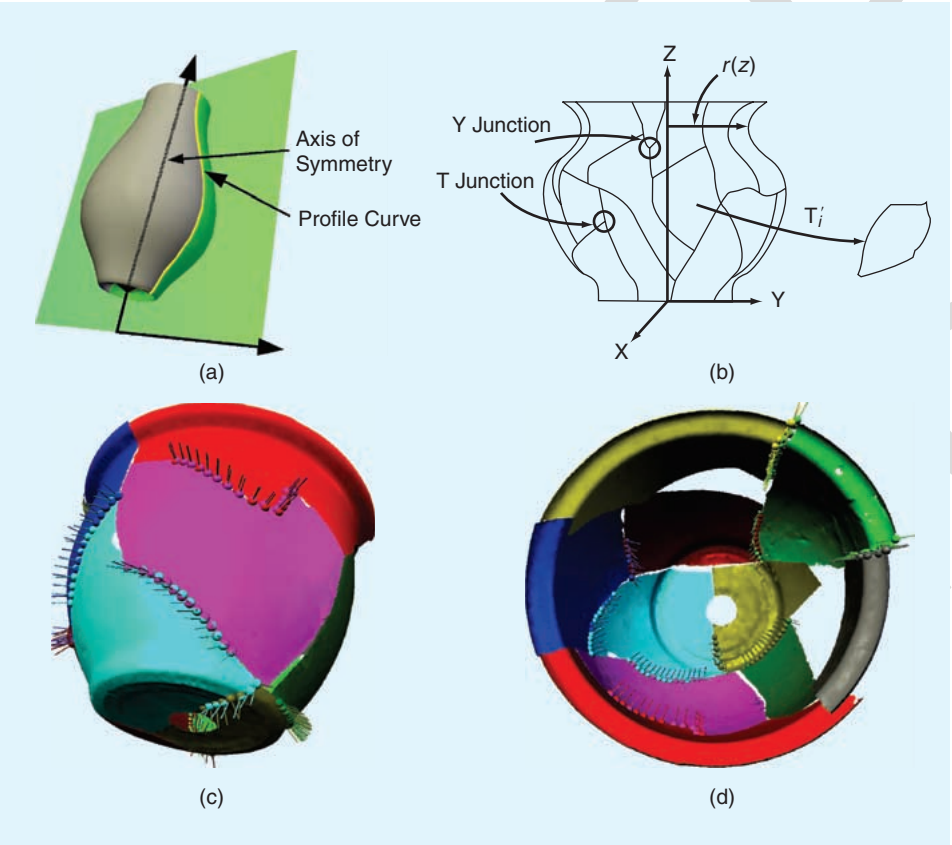

[FIG9] (a), (b) The geometric features used to model broken axially symmetric pots and (c), (d) show the result of a pot reconstruction result. (a) Shows the outer surface features for the pot, Ω_{outer}, which specify a surface of revolution as represented by a 3-D-line, the pot central axis, and a planar profile curve with respect to that the axis. (b) The outer break curve features, $\Omega_{outer\perp}$. (c) and (d) Two views of a 13-fragment pot reconstruction where ten of the pieces were provided to the system for reconstruction. Small spheres on the pot surface in (c) and (d) denote locations of matched boundaries.

In a sequence of papers [3], [16], [17], and [43], a system is described for reconstructing pots from measurements of the pot fragment outer surfaces, outer fracture curves, and their junctions. This system is unique in three ways: 1) it is the only approach that makes use of the entire fragment outer surface, Ω_{outer}, by forcing the reconstructed solution to be a surface of revolution, i.e., by combining local and global geometric information within a single performance functional to be extremized; 2) it is the only approach that uses a Bayesian framework for reconstruction; and 3) it is the only approach that addresses the problem of solving multiple puzzles simultaneously.

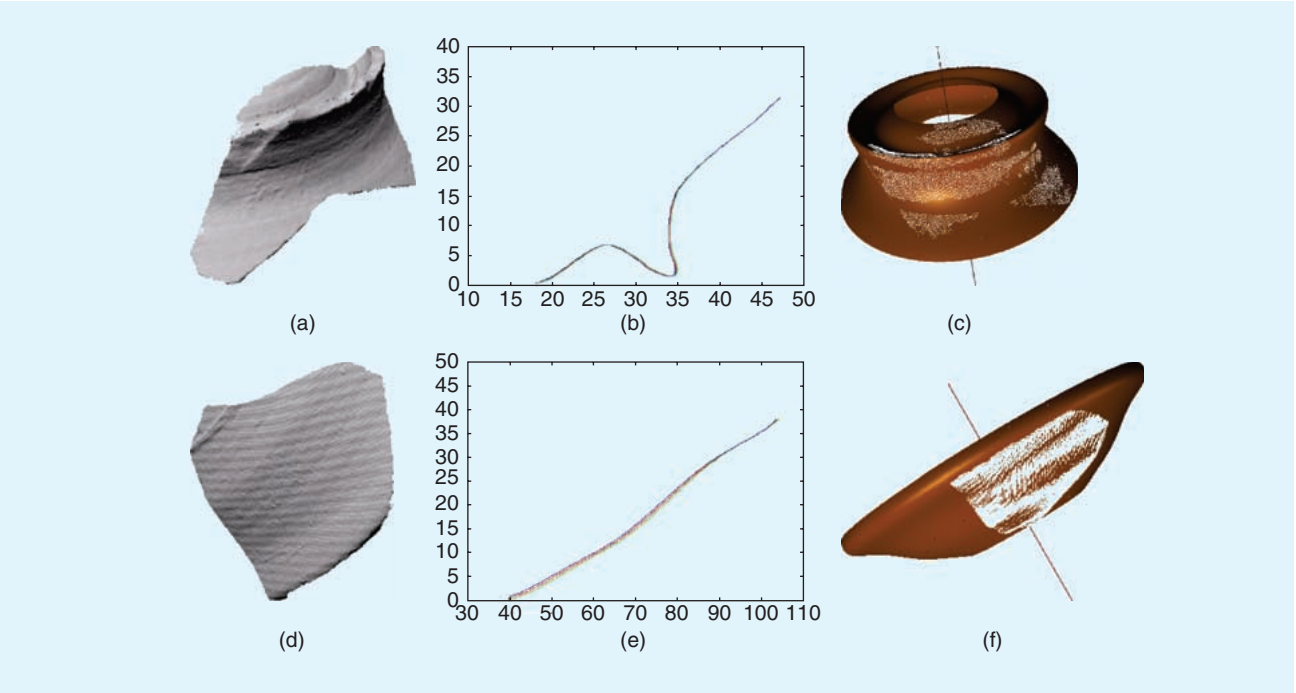

[FIG10] How geometric features of the vessel surface, i.e., the pot axis and associated profile curve, are estimated from digitized outer pot fragment surfaces. (a) (d) Two fragment outer surfaces excavated from the Great Temple in Petra, Jordan, are shown. (b) (e) The maximum likelihood estimates of the fragment outer surfaces shown in (a) and (d). The plotted curves in (b) and (e) show the fragment surface radius (x-axis) as a function of height along the pot axis (y-axis) the curve thickness indicates variability in the estimated surface as estimated by the bootstrap method, i.e., resampling the data and re-estimating the axis/profile for each fragment 100 times. In (c) and (f), the 3-D axially symmetric surfaces represented in (b) and (c) are shown n brown with the measured fragment data points superimposed in white [42]. (Used with permission from [42].)

ASSEMBLING POTTERY FROM MEASUREMENTS OF THEIR FRAGMENTS

Given the importance of pottery in understanding our past, there has been a significant research effort put forth to deal with the analysis of pottery fragments. Yet, very few approaches seek to actually reconstruct broken pots from measurements of their fragments.

An approach in [27] seeks to classify pot fragments as coming from a set of a priori known shapes provided in a typology. Typologies consist of sets of predefined prototype shapes that serve as the characteristic shape of a given vessel. The prototype shape may be only rims, bases, or the entire pottery vessel shape. While classification of pottery fragments to typologies is quite useful for archaeologists, these techniques are not designed to match actual measured fragments to estimate the shape of the unbroken pottery vessels. Hence, such techniques are seen as classification systems rather than systems for reconstruction.

Other approaches include [44], which is a volumetric approach. Here the authors seek to compute a volumetric model of a pot by extruding a spatial volume defined by revolving the pot fragments about the central pot axis. In [45] and [46], the authors use measured 3-D geometry of fragments to generate 3-D models which are then assembled manually by using a graphics interface on a computer.

FEATURE SELECTION

The features used in approach [43] consist of 1) the fragment outer surface, Ω_{outer}, as fit by a surface of revolution in terms of a central axis and a profile curve with respect to that axis, and 2) the outer break curves, $\Omega_{\text{outer}\perp}$, which, as these curves are constrained to lie on the axially symmetric surface of the pot, are similar in many ways to the fracture curve boundaries from 2-D tablet matching. $\Omega_{\text{outer}\perp}$ is modeled as sequences of 3-D points with respect to the pot axis, and their (x, y, z) coordinates constitute the variables of $\Omega_{\text{outer}\perp}$. Similar to the approach in [1], the outer break curve segments come together at junctions. As these junctions are 3-D points constrained to lie on the 2-D axially symmetric surface, we call these locations outer vertices that typically correspond to high curvature locations along the outer fragment surface break curves (see Figures 9 and 10). Note in Figure 9(b) that the pot profile curve captures the variations in the pot surface radius as a function of height along the pot radius. The profile curve is modeled as an algebraic curve of degree six, i.e., as the zero set of a sixth degree implicit polynomial. This representation effectively captures the outer surface model as shown in Figure 10, where one fragment is damaged and the other is almost flat. The axis/profile-curve model that is common to both sherds i and j when the sherds are aligned is denoted $\Omega_{\text{outer}(i,j)}$, and the part of $\Omega_{\text{outer}\perp}$ that is common to that portion of the outer break-curve along which aligned fragments i and j broke apart is denoted $\Omega_{\text{outer}\perp(i,j)}$.

MATCHING PERFORMANCE FUNCTIONAL

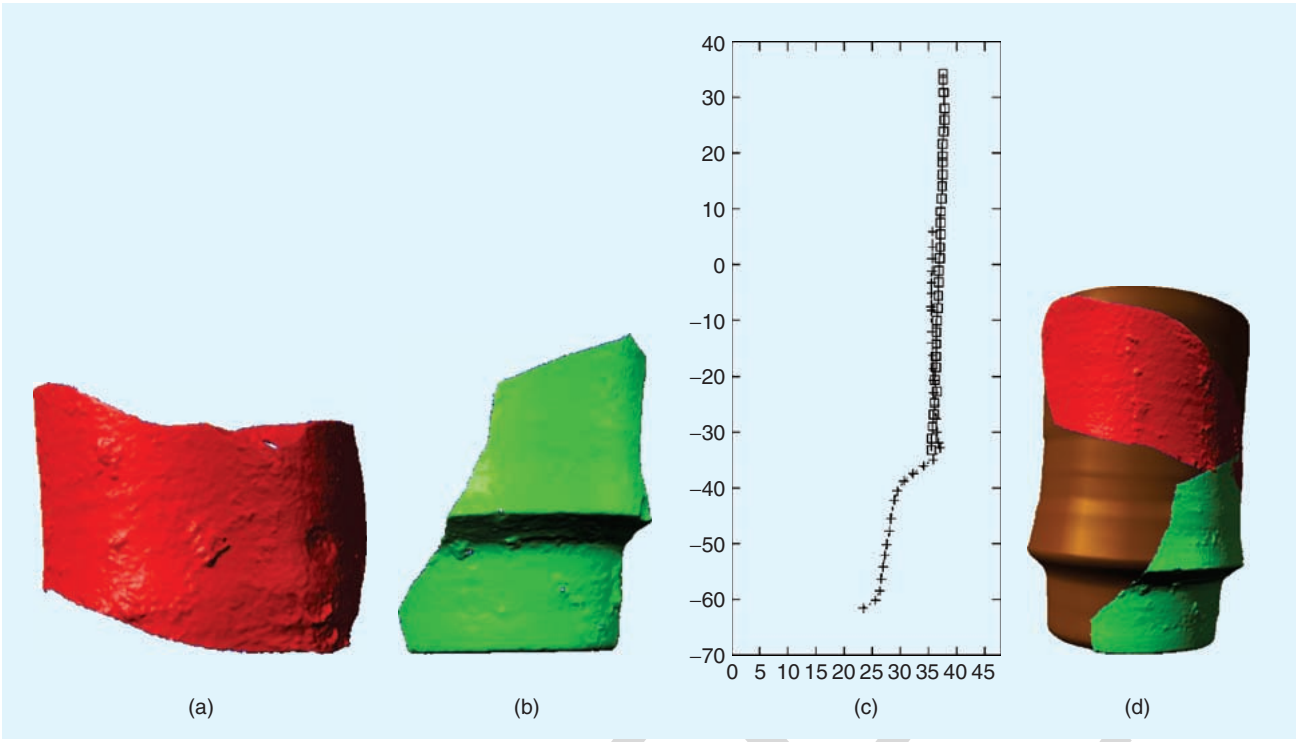

[FIG11] (a) (b) Surface measurement meshes from archaeological fragments of a Nabatean cup which are to be aligned. (c) Here the aligned mean vectors of the estimated axis/profile-curve distributions for the fragment pair are shown. (d) The 3-D axially symmetric surface generated by merging the two distributions is given. It is shown as a brown surface mesh with the aligned fragment data superimposed [16]. (Used with permission from [16].)

Because pot reconstruction in [3] and [43] is treated as Bayesian maximum a-posteriori (MAP) estimation, the matching performance functional is determined. For a pair of fragments having their data sets in some aligned position, the matching performance functional depends on the outer break curve data for each fragment and the outer surface data for each fragment. Note that only the portion of outer-break curve along which the two fragments match is used. Denote this data by $C_{(i,j)}$. Similarly, for the two outer-surface data sets, one for each fragment, only that data that is pertinent to the patch of axis/profile-curve shared by the two fragments is used. Denote this data set by $\mathcal{D}_{(i,j)}$. In words, the matching performance functional for the aligned pair of fragments, i and j, is P(break-curve data, outer-surface data | outer break-curve variables, axis/profile-curve variables), which, in equation form, is

$$P(C_{(i,j)}, \mathcal{D}_{(i,j)} | \Omega_{\text{outer}\perp(i,j)}, \Omega_{\text{outer}(i,j)}). \tag{2}$$

Equation (2) assumes the fragment data alignment is fixed, then the values used for $\Omega_{\text{outer}\perp(i,j)}$ and $\Omega_{\text{outer}(i,j)}$ are the values for which (2) is maximum, i.e., the maximum likelihood estimate (MLE) for these. In (2) it is assumed the data for the two fragments are aligned. Alignment means a Euclidean transformation $\mathbf{T}_{(i,j)}$ (3-D translation and rotation) of one of these data sets with respect to the other. Hence, $C_{(i,j)}$ and $\mathcal{D}_{(i,j)}$ depend on $\mathbf{T}_{(i,j)}$, and the conditioning should, therefore, also be on $\mathbf{T}_{(i,j)}$, and the resulting (2) should then be maximized with respect to $\mathbf{T}_{(i,j)}$, also. Then the matching performance functional for a pair of fragments is

$$P(C_{(i,j)} | \widehat{\Omega}_{\text{outer}\perp(\text{i,j})}, \widehat{\Omega}_{\text{outer(i,j)}}, \widehat{T}_{(i,j)}) P(\mathcal{D}_{(i,j)} | \widehat{\Omega}_{\text{outer(i,j)}}, \widehat{\mathbf{T}}_{(i,j)}) \tag{3}$$

where (2) factored into a product because $C_{(i,j)}$ and $\mathcal{D}_{(i,j)}$ are assumed to be conditionally independent. What are the explicit expressions for the two factors in (3)? It is assumed that a curve data point is a 2-D circularly symmetric Gaussian perturbation, within a plane, of a point on the curve. The point on the curve is that which is closest to the data point, and the plane is the plane perpendicular to the curve at the point on the curve. It is assumed that a surface data point is an independent identically distributed Gaussian perturbation, along a line perpendicular to the surface at the point on the surface closest to the data point. For example, assume a plane specified by a point in the plane, $\mathbf{p}_0$, and unit vector perpendicular to the plane, $\mathbf{w}$, then the probability density for the surface measurement point $\mathbf{p}$ is given as

$$P(\mathbf{p}|\mathbf{w}, \mathbf{p}_0) = \frac{1}{\sqrt{2\pi\sigma_s^2}} \exp\left(-\frac{1}{2\sigma_s^2}[\mathbf{w}^t(\mathbf{p}-\mathbf{p}_0)]^2\right). \tag{4}$$

(There is a reasonable rationale for these models [3].) Then each factor in (3) is a product of Gaussian probability density functions for a set of points, and (3) is a constant multiplying an exponential which has an exponent which is a sum of two quadratic forms (energy functions), one for each set of points. One energy function is a measure of how well the two outer-break-

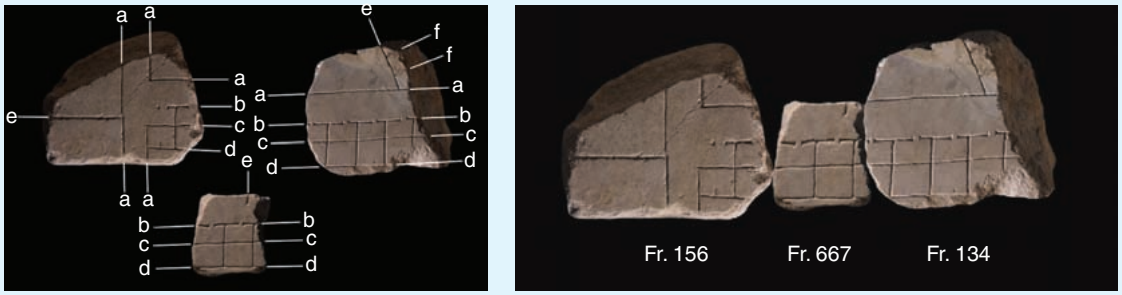

[FIG12] In one approach from [2], (a) incised fragment features on the outer surface boundary are annotated by hand and complemented by automatically computed geometric features. Features for each fragment are then exhaustively compared to features of all other fragments to identify potential fragment matches to be reviewed by an expert user. (b) A correct arrangement of the three fragments computed by the automatic matching algorithm. Note that in (b) the match between Fr. 156 and Fr. 667 was previously unknown to the archaeological community, the other match was discovered in 1992.

curve data sets, one from each curve, represent a common curve, and the other energy function is a measure of how well the two outer-surface data sets represent a common axially symmetric surface.

SEARCHING FOR THE SOLUTION

Initial approaches to this problem, attempted [43], to compute the geometry of the pot as a maximum likelihood estimation problem where we seek to find the collection of matched outer surface break curves and global axis/profile-curve geometry that maximizes the probability of all the measured fragment data. To align a pair of fragments, the authors align a pair of outer vertices (points of high curvature on the outer break-curve) one on each fragment, and then align the fragments keeping these vertices in close proximity. There are roughly four vertices per fragment and two candidate matching curves for each vertex. If there were 100 fragments, the number of alignments required to compare a pair of fragments is 64, and there are 4,950 ways to choose the candidate fragment pair from the 100 fragments, generating a total of 316,800 alignments for all possible pairwise alignments. The search algorithm starts with all of these possible configurations of two fragments, arranges them in a stack by order of decreasing probably or equivalently, increasing energy, and then grows configurations by recursively merging configurations in the stack. Initially, the fragment pairs in the stack are merged into triplets using pairs which include a common fragment. Merged configurations are reinserted into the stack according to their probability as modeled in (3). This process continued recursively, yielding larger configurations of fragments and enforcing a search rationale that considers highly probable fragment configurations first and less-likely matches later. The backtracking capabilities of this search method is limited by a set of thresholds, which limit the total number of triplets, quadruplets, etc. that may be considered part of the system. In this way, high probability configurations are grown and configurations will automatically grow for different pots. Note, configuration probabilities or energies must be appropriately normalized here because different configurations involve different numbers of data points.

Results for this approach are shown in Figure 9(c) and (d). While this method works, due to the computational complexity of the search problem, i.e., finding correct fragment matches, and the computational cost of each fragment comparison, this method can handle only small numbers of fragments. The computational killer here is that this approach processes the raw range data each time two or more surface fragments are aligned and the number of data points used for each fragment is typically in the many thousands. A recent variation on this approach, [16], uses compact probability distributions for the fragment geometry rather than repeatedly using the same raw data thus speeding up the matching process and allowing for quicker or more comprehensive searches for compatible matches. Also, rather than starting alignments by aligning break-curve vertices, it is orders of magnitude faster to first align axis/profile-curves for groups of fragments (see Figure 11). A final comment is that use of the global structure of axial symmetry provides robustness to erosion and chipping of the outer break-curves. This is because the axis/profile-curve estimates are insensitive to these deformations, and the axis/profile-curve estimates provide such powerful information for fragment reassembly that perfect fitting of outer break-curves where pot fragments come together is no longer critical.

BUILDING THE FORMA URBIS ROMAE

Members of the Forma Urbis Romae project at Stanford have attempted three very different 3-D fragment matching approaches. The cumulative efforts of these approaches have resulted in 20 new fragment matches and placements of high likelihood. It is worth noting that such efforts in assembling very difficult archaeological puzzles succeed by relying on prior knowledge and archaeological expertise to verify suggested matches, identify false positives, and take into account information that is difficult to computationally model such as evidence from excavation data and literary sources.

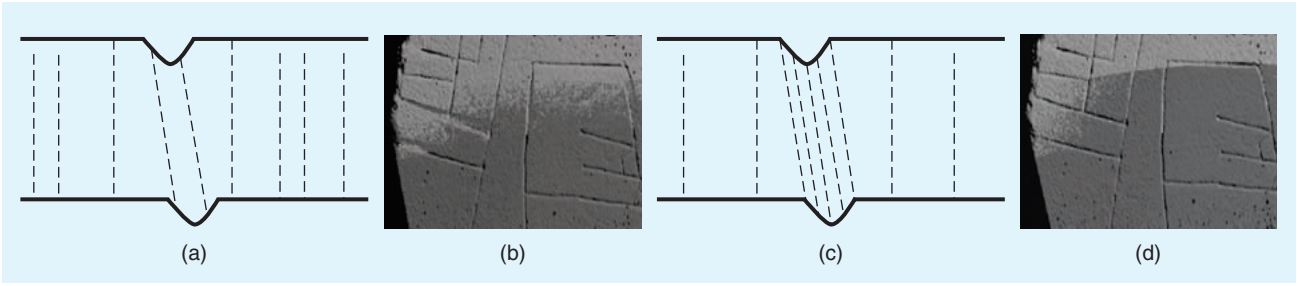

[FIG13] In [23], the authors devise techniques for extremely accurate surface scan alignments. Such techniques were necessary to extract the subtle and extremely informative locations of incised features on the tablet surface. (a) Here are two curves as solid black lines and a correspondence between the curves as dashed lines between corresponding curve samples. As indicated in (a), alignment techiques such as ICP choose these samples uniformly over the region to be aligned. (b) Shown here is the alignment obtained by minimizing the average pointwise error over this set of corresponding points. (c) The same curves with a new correspondence generated by sampling more densely in regions where the curve changes shape, i.e., where the curve normal varies. (d) This normal-space sampling strategy allows for much more accurate alignments. This is especially important for nondescript (flat) surfaces that are sparsely populated with distinctive features which is precisely the case for the tablet fragments of the Forma Urbis Romae. (Used with permission from [23].)

FEATURE SELECTION

One approach, shown in Figure 12, seeks to match incised features that extend off the boundary of the fragment outer surface, Ω_{outer}. These incisions are manually annotated to indicate relative locations of incised features, their position and direction, and the incision feature type (such as rows and columns, tabernas fronts, aqueducts, etc.).

In another approach mentioned in [2], the project team attempted to reconstruct the fragments by matching the fracture surfaces of fragments. Yet, the advanced state of deterioration of the fracture surfaces for these fragments has made assembling the fragments in this way difficult.

A third approach that is reported to be more successful uses identifiable clamp holes and masonry patches present on the wall where the map was previously hung and the complementary interlocking geometry found on the fragments to constrain the possible locations for fragments that include such features.

MATCHING PERFORMANCE FUNCTIONAL

Much emphasis within this project has been placed in obtaining accurate alignments of the measured range images. In [23], Rusinkiewicz and Levoy evaluate many variations of the pervasive iterative closest point (ICP) algorithm initially proposed in [21] and [47], which is commonly used to compute the Euclidean alignment between two overlapping 3-D surfaces. ICP algorithms typically select a random subset of all the available points for estimating an alignment. However, for the flat tablet surfaces, extra emphasis must be placed on having extremely precise alignment of incised surface features to produce accurate range image alignments. Towards this end, the authors suggest a new selection criterion for the set of alignment points. Rather than choosing a random subset of the surface points uniformly distributed over the surface, they choose the set of surface points for which the normals at these points provide the widest variation in normal directions. As shown in Figure 13, this "normal-spaced" sampling of the surface points generates a nonuniform sampling of the surface that samples curved surface regions more densely than in flat surface regions. As shown in Figure 13, such considerations are necessary to accurately align pieces that are mostly flat with small subtle incised features as is the case for fragments from the Forma Urbis Romae. Given this careful technique for alignment, features such as locations where incised markings touch the fragment outer surface boundary may be known with high precision. The matching functional is then simply the sum of squared differences between these computed feature values.

SEARCHING FOR THE SOLUTION

For each of the three methods, the authors apply a straightforward brute force search guaranteed to find the global minimum since computational cost is not a major concern. This search technique exhaustively matches every computed feature vector against all other feature vectors for all other fragments generating a list of match scores ordered with the best-fit matches at the top. The best matches can then be manually reviewed for correct matches such as the triplet of fragments shown in Figure 12(b).

GENERIC 3-D OBJECT RECONSTRUCTION

We now discuss approaches proposed in [9] and [11] for the reconstruction of generic 3-D free-form surfaces. Both approaches follow the generic assembly process outlined in Figure 4 but differ in the following three aspects: i) the features selected for matching, ii) the matching performance functional, and iii) how correct fragment-pair matches are detected and merged into larger puzzle solutions. Both approaches also discuss the problem of identifying the outer fragment surface, Ω_{outer}, and the fracture surface, Ω_{inner}, automatically from 3-D scan data. We discuss the difference between the two approaches in the following sections.

FEATURE SELECTION

In [9], the authors use scanned data points of the fracture surface, Ω_{inner}, and a curvature and torsion representation of the exterior fracture space curves, $\Omega_{\text{outer}\perp}$, to reconstruct 3-D objects. Features are extracted automatically from the fracture

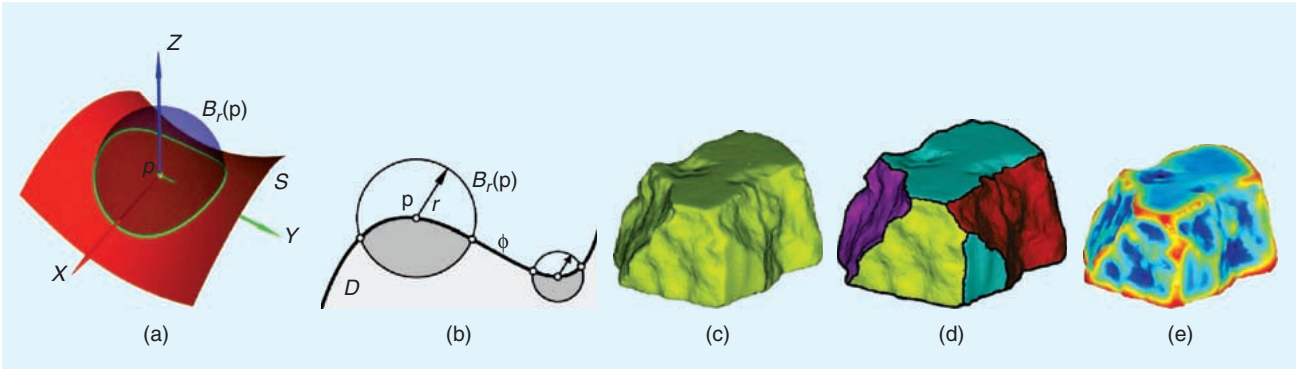

[FIG14] These images demonstrate how the authors compute integral invariants (a) and (b) [see (5) and (6)] and associated text for details). (c) and (e) show how a raw digitized fragment (c) is segmented, (d) compactly represented in terms of one exemplar invariant feature: (e) surface sharpness. These invariant features are further compacted and then used for fracture surface and break curve matching. (Used with permission from [11].)

surface using a heuristic region-growing algorithm that seeds a single random polygon on the surface mesh. The surface region is grown by adding those polygons along the boundary of the region, i.e., polygons that share an edge with the region, that have surface normals roughly parallel to those included within the region (some variability is permitted). This process continues until all potential neighbors for a given region have been included. At this point, if the surface is not covered with regions, a new region is started by randomly selecting a polygon not associated with any region. This process continues until the surface is completely covered with small regions. Adjacent regions with similar orientations are subsequently merged to generate the complete set of fragment surface regions. Each of the resulting regions is then inspected for roughness by computing the spatial variability of the surface points within the region. Those that have variability higher than a threshold are designated as fracture surfaces. As in 2-D, a multiscale curve matching process for interior fracture boundaries is applied to quickly eliminate matches that are clearly incorrect (see the section "Matching Performance Functional" for details). This process low-pass filters the fragment outer break curves with a radially symmetric Gaussian kernel to allow for quick matching and reducing the influence of noise on computed curvature and torsion values for this space-curve similar to the process shown in Figure 6.

In [11], the authors use integral invariant features of fracture surface patches, Ω_{inner}, and break curves, $\Omega_{\text{inner}\perp}$, for reconstruction. As shown in Figure 14(a) and (b), these features are computed at a surface point by integrating spatial functions over the region contained within a sphere of radius r centered at the surface point. One of several features used in their approach is a volume descriptor defined in

$$V^r(\mathbf{p}) = \frac{3}{4\pi r^3} \int_{B_r(\mathbf{p})} \chi_D(\mathbf{x})\, d\mathbf{x}, \tag{5}$$

where $\mathbf{p}$ is the 3-D surface point where we wish to compute the feature, $B_r(\mathbf{p})$ is a sphere of radius r at $\mathbf{p}$, and χ_D is the indicator function for the interior of the object, i.e., $\chi_D(\mathbf{x}) = 1$ for all $\mathbf{x}$ inside the surface and is otherwise zero [see Figure 14(b)]. This feature is functionally related to the mean surface curvature, H, by the following formula:

$$V^r(\mathbf{p}) = \frac{1}{2} - \frac{3}{16} H \cdot r + O(r^2). \tag{6}$$

As in [9], efforts are made to separate the outer surface and the fracture surface of each fragment. This is accomplished by having the user manually indicate points that lie on outer surfaces and fracture surfaces to train a two-class classifier capable of discriminating between fragment outer and fracture surfaces. The classifier is then applied to classify all of the measured surface data. The final segmentation of a fragment involves not only identification of the outer and fracture surface but also the subdivision of the fracture surface into fracture surface patches that are delineated by 3-D break curves, $\Omega_{\text{inner}\perp}$. This is accomplished by first computing a graph $G(F, E)$ where each graph node is a polygon and each graph edge is an edge shared by two mesh polygons. Each graph edge has two costs: 1) the difference in neighboring polygon surface roughness descriptors [see (7a)] and 2) the sum of neighboring polygon surface sharpness descriptors [see (7b)], which is similar in value to the norm of the surface mean curvature [see (6)]

$$e_r(\mathbf{p}) = \frac{1}{k} \sum_{i=1}^{k} \frac{\|\mathbf{n}_p - \mathbf{n}_{q_i}\|^2}{\|\mathbf{p} - \mathbf{q}_i\|} \tag{7a}$$

$$s_{\text{vol}}(\mathbf{p}) = \left(V^r(\mathbf{p}) - \frac{1}{2} \right)^2, \tag{7b}$$

where $\mathbf{p}$ is the surface point of consideration, $\mathbf{q}_i$ is one of k surface points within a radius r from $\mathbf{p}$, $\mathbf{n}_p$ is the surface normal at $\mathbf{p}$, and $\mathbf{n}_{q_i}$ is the surface normal at $\mathbf{q}_i$. In practice, the authors compute these descriptors at varying scales by changing r, the radius of the integrating sphere, and summing the computed values.

The final segmentation of the fracture surface into fracture surface patches bounded by break curves is obtained by repeatedly applying several variations on the graph cut algorithm (or normalized graph cut algorithm) on the graph $G(F, E)$ where

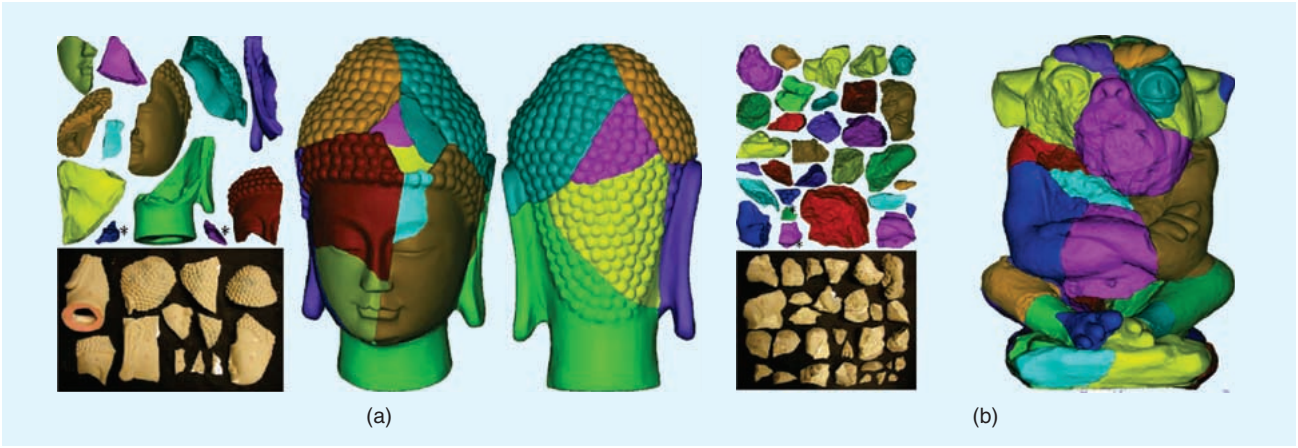

[FIG15] Results from [11] show the reconstruction of two statues broken for the purposes of experimentation. In both (a) and (b) the original fragments are shown laid out on a table (bottom left) and the digitized fragments are shown (top-left). Views of the final reconstructions are shown on the right. In (a) there were 12 fragments available, ten of which were correctly reassembled in 1.15 minutes. In (b) there were 30 fragments, and 28 of them were correctly reassembled in 3.54 minutes. (Used with permission from [11].)

soft constraints are imposed to favor selection of closed break curves that have long arc length [see Figure 14(d)].

MATCHING PERFORMANCE FUNCTIONAL

The matching metric applied in [9] and [48] for 3-D curve matching follows that in [26], which uses a similarity matrix whose contents at (row, column) index (i, j) is the sum of squared differences between the curvature and torsion at sample i from curve C and the curvature and torsion at sample j from curve $\overline{C}$. Similar segments of boundary curve are then detected as diagonal sequences of small values within the similarity matrix, i.e., there is no elastic matching in this approach. After using the similarity matrix to determine correspondences between fragment outer break curves, $\Omega_{\text{outer}\perp}$, the fracture surfaces, Ω_{inner}, adjacent to these boundaries are matched. This is accomplished by using the similarity matrix correspondence to align the matched outer break curves using [40]. Afterwards, fracture surfaces are matched by representing each of the fracture surfaces as an explicit function of the plane having normal equal to the average fracture surface normal resulting in a depth map $z = f(x, y)$. Matching surfaces is then accomplished by matching one depth map to another after inverting the z-axis of the second depth map (which is equivalent to reflecting the fragment fracture surface) so that it may be joined face to face with the other fracture surface.

In [11] fragment matching is done by clustering similar integral invariant features for different fracture surface patches, Ω_{inner}, and their adjacent break curves, $\Omega_{\text{inner}\perp}$, over a sequence of coarse-to-fine scales. The authors start by computing several invariant descriptors defined at different scales for each fracture surface patch, i.e., different values for r in (7), resulting in a vector of integral invariant features for each fragment fracture surface point. Each of the observed values for each feature is then quantized into k levels (typically 32) that collect computed features into k subgroups based on their quantized values. Within each of these subgroups, clusters are formed. Each cluster includes only those features that lie on a single surface patch and whose associated points are connected by edges on the surface mesh. Since each cluster consists of Euclidean-invariant features, i.e., integral invariants, the clusters may be directly matched to identify potential matches between small surface patches. Before doing so, each cluster is augmented with additional (Euclidean-variant) parameters. These parameters include the barycenter of the surface points associated with each cluster and a collection of representative surface points within a fixed distance of the barycenter. Pairwise matching is then accomplished by performing principal component analysis (PCA) on the augmented feature vector. Clusters which have similar size and eccentricity as represented by the sum and ratio, respectively, of the eigenvectors of the cluster scatter matrix are considered to be candidate matches. Hence, if two clusters are similar in shape as defined by their PCA decomposition and pass an additional topological verification step, they are considered a potential match and are passed through a second, more rigorous, matching criterion. The second matching criterion checks geometric consistency by ensuring that distances between corresponding points on corresponding surfaces have approximately the same point-to-point distance. Given this test is also passed, the surface points are aligned using [40] since the correspondences between the surface points has already been determined by the matching criterion applied in feature space. These pairwise fragment matches are assigned a match quality score that is the total surface sharpness of the matched region minus the combined alignment error for the matched fracture surface patches and break curves.

SEARCHING FOR THE SOLUTION

Reference [9] applies a brute force search first for matching $\Omega_{\text{outer}\perp}$ the fragment outer break curves, and then a second brute force search is applied on the surfaces adjacent to well-matched portions of the outer break curve. The best matches from this surface matching procedure is the overall puzzle solu-

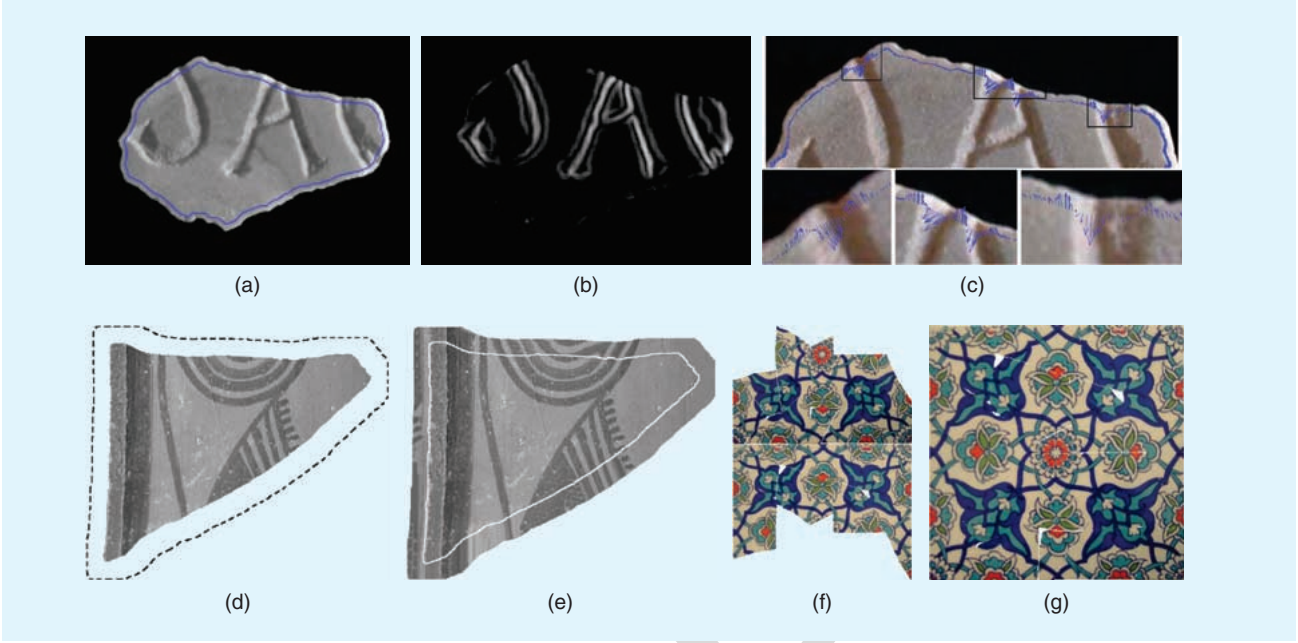

[FIG16] (a), (b), and (c), taken from [1] demonstrate one technique for matching fragments using patterns apparent on the fragment outer surface. This technique matches three features from the fragment surface along a slightly inset boundary [shown as the blue curve in (a)]. The features are (a) pixel luminance values along the boundary, (b) the gradient magnitude along the boundary, and (c) the energy in a continuation curve that span the fracture in the direction of the image gradient vectors. (d), (e), (f), and (g), taken from [13], show a second approach for appearance-based fragment matching that uses texture in-painting to extend the surface pattern (e) on the piece (d) to predict the patterns that appear on matching pieces (see text for details). (Used with permission from [1] and [13].)

tion.

The search for a global solution in [11] is accomplished by defining a graph where each fragment is a node in the graph and each edge is a candidate matching fragment detected during the pairwise matching stage. Groups of edges in this graph correspond to collections of aligned fragments, and the task is to select the appropriate subgraph associated with the correct puzzle solution. Towards this end, the authors proceed by selecting a group of edges from the graph and then merge these edges into a new subgraph. They then validate the match implied by the newly generated subgraph by making sure that the matched fragments to not penetrate each other. If the merge passes this test, it is used as a new node in the graph which may then be iteratively merged with other unmatched fragments. Presumably, this process continues until a single node exists corresponding to the global puzzle solution. Figure 15(a)–(b) shows two figurines broken for the purpose of experimentation and reassembled using technique [11].

BEYOND GEOMETRY: USING PATTERNS ON THE FRAGMENT SURFACE

In [1], [13], and [14], the authors include information obtained from patterns on the outer fragment surface to augment the matching performance functional. Figure 16(a)–(c) shows how patterns apparent on an inscribed Roman marble tablet discovered in Petra, Jordan, are used for improved fracture curve matching in [1]. In Figure 16(a), the author computes the intensity profile of the fracture curve along a curve inset a small distance from the boundary (shown in blue). The matching performance functional used in this case is a combination of the difference in matched pixel luminances [Figure 16(a)], the difference in matched gradient magnitudes [Figure 16(b)], and a curve continuation metric called the Euler spiral completion curve that seeks to complete curves across the fracture in the directions of the gradient vectors shown in Figure 16(c) while minimizing the curvature variation of the estimated completion curve.

In [13], and [14], the authors use in-painting and texture synthesis methods to extend the patterns on the outer surface of fragments to predict those patterns that appear on matching pieces [see Figure 16(d) and (e)]. The mean and variance of the predicted pixel values within small windows around the image boundary are then computed as the fragment features. Note that these are intrinsic features of the fragment, like curvature, and are also Euclidean invariant. These features are then matched with those of other fragments by transforming one piece with respect to the other and subsequently taking a weighted Euclidean distance between features at corresponding locations where the weighting coefficients decrease for the matched features close to the boundary of the extended region, i.e., matches closest to the original unextended fragment boundary have the highest weights. This matching procedure is then a weighted cross-correlation of the image features along the extended boundary and must be optimized with respect to the variables of the unknown transformation matching the fragments. The fast Fourier transform is then used to efficiently

solve for the translation variables of the transformation followed by an optimization on the unknown one-parameter rotation to determine the best relative alignment of the fragments. The search procedure is a best-first approach that searches through the fragment matches exhaustively to compute a global solution which may change depending upon the chosen order of fragments to match as shown in Figure 16(f) and (g).

DISCUSSION AND FUTURE WORK

Of the 2-D reconstruction systems discussed, correct or nearly correct results for large (~22 piece) synthetic puzzles have been reported in several cases [1], [10], [49]. Yet only [1] and [49] provide results for archaeological fragments and these include only small groups of matched fragments, i.e., pairs. Of the 3-D reconstruction systems discussed, [2], [11], and [16] all present reconstructions of artifacts although [2] is not automatic and performs reconstruction with a large amount of outside help from expert users. Reconstructions including four archaeological pot fragments from the Great Temple at Petra site are presented in [16] and three fragments from the Forma Urbis Romae project are presented in [11]. While this may initially seem discouraging, one must remember that one major contributing factor for such small numbers of matched archaeological data is, in part, due to difficulties in quickly and easily collecting and digitizing the fragments. At present, practical reassembly requires beginning with a collection of fragments that have been excavated near one another so that the likelihood of many pieces having come from the same object is high. Development of quick and accurate scanning technologies already under development promise to streamline the reassembly process making these computational aids applicable in real-world archaeological settings.

As one can see from this survey, a wide variety of approaches for artifact reconstruction have been investigated using many conceivable pieces of geometric or appearance information. Yet, we note that much information available from the outer surface remains unused in most current approaches. References [1] [50] are examples that use such information by incorporating tangent contunuity constraints across matched outer surface boundaries to construct matches that are both continuous and smooth across the joined outer surface boundary. Yet, with the exception of [43], no current methods incorporate higher order geometric features in the vicinity of the outer surface boundary. Since artifact outer surfaces are typically smooth, we feel that these features may be computed accurately and be informative enough for fragment matching to improve the current state-of-the-art. Use of this information is quite applicable to both free-form 3-D reconstruction and 2-D tablet reconstruction problems. As noted in the third approach explored by the Forma Urbis Romae project, being able to identify the outer boundary of fragments constrains the potential matches for these fragments considerably, which is a big performance boost when searching for potential candidate matching fragments. This is an area that will likely receive more consideration as scientists develop the next generation of artifact reconstruction systems.

Another aspect to this problem in need of attention is the development of an information-theoretic basis for searching compatible matches. Current methods either have an extremely difficult search problem with many false positives [1], [50] or sufficiently discriminative features to quickly find the solution with a simple matching procedure [11]. Making expeditious use of computed features requires knowledge of how much discriminative information a given feature contains, a central concept in machine learning. Seminal work in this regard is offered in [32] where the authors analyze the information content of 2-D ceramic tile fracture boundaries in terms of their curvature signatures. Using the discrete sine transform, the authors propose a method for measuring the amount of information present in a 2-D boundary curve. Search procedures that concentrate on high-information features identified using techniques such as this also show promise for improving the speed, accuracy, and reliability of reconstruction systems. A related aspect of the problem is the development of a generative model for cracks and fractures as proposed in [51] or for erosion. Such work may lead to material-specific fracture boundary models capable of succinctly representing and robustly estimating fracture boundaries and fracture boundary junctions, i.e., the vertices and break curve junctions that provide crucial matching features and performance speed-ups to some computational reconstruction systems.

While approaches to date are sophisticated and computationally efficient, none of the published works claim to be ready for deployment in the field as a tool for archaeologists. In fact, it is surprising to note that the Forma Urbis Romae project using brute force methods and expert users has been very successful in discovering 20 new matches for a reconstruction problem that humans have been studying for hundreds of years. This encourages the use of current reconstruction systems in a cybernetic context. Such systems would use computerized artifact reconstruction to complement manual reconstruction, exploiting the powers of each to generate a system that is more efficient and accurate than either could be independently. Here, the computer can quickly identify a small set of candidate matches that the human user can manually prune to aid the computer in obtaining the correct global solution. This is an intent of [49].

Achievable goals for the next generation of archaeology fragment assembly systems are: 50-125 fragment ceramic vessels; glass vessels, and large architectural structures such as castles, palaces, temples, and habitats. Solutions to these probems will also constitute very significant contributions to digital signal processing and automatic/interactive machine learning generally.

CONCLUSION

This article has provided a survey of the state-of-the art in both automatic and semi-automatic artifact reconstruction systems. Discussion has cast the puzzle reconstruction problem into a generic terminology that is formalized appropriately for the 2-D

and 3-D artifact reconstruction problems. Two leading approaches for 2-D tablet reconstruction and four leading approaches for 3-D object reconstruction have been discussed in detail, including partial or complete descriptions for the numerous algorithms upon which these systems rely. Several extensions to the geometric matching problem that use patterns apparent on the fragment outer surface were also discussed that generalize the problem beyond that of matching strictly geometry. The models needed for solving these problems are new and challenging, and most involve 3-D that is largely unexplored by the signal processing community. This work is highly relevant to the new 3-D signal processing that is looming on the horizon for tele-immersion. We thank the authors of [1], [2], [11], [13], [23], and [28] for granting permission to use their figures in this article.

ACKNOWLEDGMENT

This work was partially supported by NSFITR Grant 0205477.

AUTHORS

Andrew Willis (arwillis@uncc.edu) is an assistant professor in the Department of Electrical and Computer Engineering at the University of North Carolina at Charlotte (UNCC) where he supervises the UNC Charlotte Machine Vision Laboratory. He received the B.Sc. in computer science and B.Sc. in electrical engineering from Worcester Polytechnic Institute in 1995. From 1995 to 1998, he worked for Morgan Construction Co. as a field service engineer. He obtained an Sc.M. in applied mathematics and an Sc.M. in electrical engineering, ultimately completing his Ph.D. in engineering sciences in 2004, all from Brown University. He is a Senior Member of the IEEE and a member of the ACM, AAAS, and ASEE. His research focuses on 3-D computer vision, computer graphics, and stochastic inference for geometric problems.

David B. Cooper (cooper@lems.brown.edu) received the B.Sc. and Sc.M. degrees in electrical engineering from MIT in 1957 and the Ph.D. degree in applied mathematics from Columbia University, New York, in 1966. He worked for Sylvania Co. and then the Raytheon Co. from 1957–1966. Since 1966, he has been a professor of engineering at Brown University, Providence, Rhode Island. His more recent research interests have been on 2-D and 3-D object estimation and recognition from many unreliable fragments of information obtained from video or LADAR, geometric learning, adhoc sensor networks–specifically the theory for how to use 1,000 cameras–and various data analysis problems for extraction of information from fragments found at archaeological sites. He is an IEEE Fellow.

REFERENCES

[1] J. McBride, "Archaeological fragment re-assembly using curve matching," M.S. thesis, Brown Univ., Div. Eng., Providence, RI, Sept. 2002.

[2] D. Koller, J. Trimble, T. Najbjerg, N. Gelfand, and M. Levoy, "Fragments of the city: Stanford's digital forma urbis romae project," *J. Roman Archaeology Suppl.*, vol. 61, pp. 237–252, Mar. 2006.

[3] A. Willis, D. Cooper, S. Andrews, J. Baker, Y. Cao, D. Han, K. Kang, W. Kong, F.F. Leymarie, X. Orriols, S. Velipsalar, E.L. Vote, M.S. Joukowsky, B.B. Kimia, D.H. Laidlaw, and D. Mumford, "Assembling virtual pots from 3-D measurements of their fragments," in *Proc. VAST Int. Symp.Virtual Reality Archaeology and Cultural Heritage*, 2001, pp. 241–253.

[4] H. Freeman and L. Garder, "A pictorial jigsaw puzzles: The computer solution of a problem in pattern recognition," *IEEE Trans. Electron. Comput.*, vol. 13, no. 2, pp. 118–127, 1964.

[5] H. Wolfson, E. Schonberg, A. Kalvin, and Y. Lambdan, "Solving jigsaw puzzles using computer vision," *Ann. Oper. Res.*, vol. 12, pp. 51–64, Feb. 1988.

[6] G. Radack and N. Badler, "Jigsaw puzzle matching using a boundary-centered polar encoding," *Comput. Graph. Image Process.*, vol. 19, pp. 1–17, May 1982.

[7] D. Goldberg, C. Malon, and M. Bern, "A global approach to automatic solution of jigsaw puzzles," in *Proc. Conf. Computational Geometry*, 2002, pp. 82–87.

[8] S. Li, *Markov Random Field Modeling in Image Analysis*. New York: Springer-Verlag, 2001, ch. 7.

[9] G. Papaioannou, E.-A. Karabassi, and T. Theoharis, "Reconstruction of three-dimensional objects through matching of their parts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 1, pp. 114–121, 2002.

[10] J. Stolfi and H. Leitão, "A multiscale method for the reassembly of two-dimensional fragmented objects," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 9, pp. 1239–1251, 2002.

[11] Q.-X. Huang, S. Flöy, N. Gelfand, M. Hofer, and H. Pottmann, "Reassembling fractured objects by geometric matching," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 569–578, 2006.

[12] W. Kong and B.B. Kimia, "On solving 2-D and 3-D puzzles using curve matching," in *Proc. Conf. Computer Vision Pattern Recognition (CVPR)*, Hawaii, Dec. 2001, pp. II-583–II-590.

[13] M. c. Sağiroğlu and A. Erçil, "A texture based matching approach for automated assembly of puzzles," in *Proc. 8th Int. Conf. Pattern Recognition*, vol. 3, 2006, pp. 1036–1041.

[14] M. c. Sağiroğlu and A. Erçil, "A texture based approach to reconstruction of archaeological finds," in *Proc. 6th Int. Symp. Virtual Reality, Archaeology, and Cultural Heritage*, 2005, pp. 137–142.

[15] B. Kimia and J. McBride, "Archaeological fragment re-assembly using curve matching," in *IEEE Workshop Computer Vision Applications in Archaeology (CVPR)*, 2003, p. 3.

[16] A. Willis and D. Cooper, "Estimating a-priori unknown 3-D axially symmetric surfaces from noisy measurements of their fragments," in *Proc. 3rd Int. Symp. 3-D Data Processing, Visualization, and Transmission*, 2006, pp. 334–341.

[17] A. Willis, Stochastic 3-D geometric models for classification, deformation, and estimation, Ph.D. dissertation, Div. Eng., Brown Univ., Providence, RI, May 2004.

[18] Y. Ma, S. Soatto, J. Kosecka, and S.S. Sastry, *An Invitation to 3-D Vision*. New York: Springer-Verlag, 2003.

[19] T. Malzbender, B. Wilburn, D. Gelb, and B. Ambrisco, "Surface enhancement using real-time photometric stereo and reflectance transformation," in *Proc. Eurographics Symp. Rendering*, June 2006, pp. 245–248.

[20] S. Rusinkiewicz, O.A. Hall-Holt, and M. Levoy, "Real-time 3-D model acquisition.," in *Proc. Int. Conf. Computer Graphics and Interactive Techniques* (*SIGGRAPH)*, 2002, pp. 438–446.

[21] P. Besl and N. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.

[22] A. Fitzgibbon, "Robust Registration of 2-D and 3-D Point Sets," in *Proc. British Machine Vision Conf.*, 2001, pp. 1145–1153.

[23] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP Algorithm," in *Proc. Intl. Conf. 3-D Digital Imaging and Modeling*, 2001, pp. 145–152.

[24] G. Turk and M. Levoy, "Zippered polygon meshes from range images," in *Proc. Int. Conf. Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1994, pp. 311–318.

[25] H.W. Guggenheimer, *Differential Geometry*. New York: Dover, 1977.

[26] G. Uçoluk and I.H. Toroslu, "Automatic reconstruction of broken 3-D surface objects," *Comput. Graph.*, vol. 23, no. 4, pp. 573–582, 1999.

[27] A. Gilboa, A. Karasik, I. Sharon, and U. Smilansky, "Towards computerized typology and classification of ceramics," *J. Archaeological Sci.*, vol. 31, no. 6, pp. 681–694, 2004.

[28] W. Kong, "On solving 2-D and 3-D puzzles using curve matching," M.S. thesis, Brown Univ., Div. Eng., Providence, RI, May 2002.

[29] E. Justino, L.S. Oliveira, and C. Freitas, "Reconstructing shredded documents through feature matching," *Forensic Sci. Int.*, vol. 160, pp. 140–147, July 2006.

[30] H.J. Wolfson, "On curve matching," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 5, pp. 483–489, 1990.

[31] M.M. Lipschutz, *Schaum's Outline of Theory and Problems of Differential Geometry*. New York: McGraw-Hill, 1969.

[32] J. Stolfi and H. Leitão, "Measuring the information content of fracture lines," *Int. J. Comput. Vis.*, vol. 65, no. 3, pp. 163–174, 2005.

[33] S. Hermann and R. Klette, "A comparative study on 2-D curvature estimators," in *Proc. Int. Conf. Computing: Theory and Applications (ICCTA)*, 2007, pp. 584–589.

[TABLE 1] VARIABLE DEFINITIONS AND TERMINOLOGY USED IN THIS ARTICLE FOR 2-D AND 3-D ARTIFACT RECONSTRUCTION (SEE FIGURE 3).

		2-D RECONSTRUCTION		3-D RECONSTRUCTION	
GEOMETRIC VARIABLE SET	**GENERIC TERM**	**GEOMETRIC PRIMITIVE**	**REFERRED TO AS:**	**GEOMETRIC PRIMITIVE**	**REFERRED TO AS:**
Ω_{outer}	OUTER BOUNDARY	PLANAR CURVE	OUTER CURVE	SURFACE PATCH	OUTER SURFACE
Ω_{inner}	FRACTURE BOUNDARY	PLANAR CURVE	FRACTURE CURVE	SURFACE PATCH	FRACTURE SURFACE
$\Omega_{inner\perp}$	FRACTURE JUNCTION	2-D POINT	VERTEX	3-D SPACE-CURVE	BREAK CURVE
$\Omega_{outer\perp}$	OUTER FRACTURE JUNCTION	2-D POINT	OUTER VERTEX	3-D SPACE-CURVE	OUTER BREAK CURVE

[34] S. Manay, B.-W. Hong, D. Cremers, A.J. Yezzi, and S. Soatto, "Integral invariants for shape matching," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 10, pp. 1602–1618, 2006.

[35] X. Tang, "A sampling framework for accurate curvature estimation in discrete surfaces," *IEEE Trans. Visual. Comput. Graphics*, vol. 11, no. 5, pp. 573–583, 2005.

[36] S. Rusinkiewicz, "Estimating curvatures and their derivatives on triangle meshes," in *Proc. 3-D Data Processing, Visualization, and Transmission, 2nd Int. Symp.*, (*3-DPVT)* 2004, pp. 486–493.

[37] S. Petitjean, "A survey of methods for recovering quadrics in triangle meshes," *ACM Comput. Surv.*, vol. 34, no. 2, pp. 211–262, 2002.

[38] M. Desbrun, M. Meyer, P. Schöder, and A.H. Barr, "Implicit fairing of irregular meshes using diffusion and curvature flow," in *Proc. 26th Annu.Conf. Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1999, pp. 317–324.

[39] G. Taubin, "A signal processing approach to fair surface design," in *Proc. 22nd Annu. Conf. Computer Graphics and Interactive Techniques SIGGRAPH*, 1995, pp. 351–358.

[40] B. Horn, "Closed-form solution of absolute orientation using unit quaternions," *J. Opt. Soc. Amer.*, vol. 4, no. 4, pp. 629–642, 1987.

[41] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 4, pp. 376–380, 1991.

[42] A. Willis, X. Orriols, and D. Cooper, "Accurately estimating Sherd 3-D surface geometry with application to pot reconstruction," in *IEEE Workshop Applications Computer Vision in Archaeology (CVPR)*, June 2003, p. 5.

[43] A. Willis and D. Cooper, "Bayesian assembly of 3-D axially symmetric shapes from fragments," in *Proc. Conf. Computer Vision and Pattern Recognition*, vol. I, 2004, pp. 82–89.

[44] M. Kampel and R. Sablatnig, "Virtual reconstruction of broken and unbroken pottery," in *Proc. Conf. 3-D Digital Imaging and Modeling*, 2003, pp. 318–325.

[45] I. Marie and H. Qasrawi, "Virtual assembly of pottery fragments using moire surface profile measurements," *J. Archaeological Sci.*, vol. 32, no. 10, pp. 1527–1533, 2005.

[46] Y. Lu, H. Gardner, H. Jin, N. Liu, R. Hawkins, and I. Farrington, "Interactive reconstruction of archaeological fragments in a collaborative environment," in *Proc. 9th Biennial Conf. Australian Pattern Recognition Society Digital Image Computing Techniques and Applications (DICTA)*, 2007, pp. 23–29.

[47] Y. Chen and G. Medioni, "Object modeling by registration of multiple range images," in *Proc. IEEE Conf. Robotics and Automation*, vol. 3, 1991, pp. 2724–2729.

[48] G. Papaioannou and E.-A. Karabassi, "On the automatic assemblage of arbitrary broken solid artefacts," *Image Vis. Comput.*, vol. 21, no. 5, pp. 401–412, 2003.

[49] C. Aras, Hindsite: A robust system for archaeological fragment re-assembly," M.S thesis, Div. Eng., Brown Univ., Providence, RI, May 2007.

[50] A. Willis, D. Cooper, S. Andrews, J. Baker, Y. Cao, D. Han, K. Kang, W. Kong, F.F. Leymarie, X. Orriols, S. Velipsalar, E.L. Vote, M.S. Joukowsky, B.B. Kimia, D.H. Laidlaw, and D. Mumford, "Bayesian pot-assembly from fragments as problems in perceptual-grouping and geometric-learning," in *Proc. Int. Conf. Pattern Recognition (ICPR)*, vol. III, 2002, pp. 297–302.

[51] B. Desbenoit, E. Galin, and S. Akkouche, "Modeling cracks and fractures," *Vis. Comput.*, vol. 21, no. 8, pp. 717–726, 2005.

[52] D. Crispell, D. Lanman, P.G. Sibley, Y. Zhao, and G. Taubin, "Beyond silhouettes: Surface reconstruction using multi-flash photography," in *Proc. 3rd Int. Symp. 3-D Data, Processing, Visualization and Transmission (3DPVT)*, June 2006.

[53] S. Liu, K. Kang, J.P. Tarel, and D.B. Cooper, "Free-form object reconstruction from silhouettes, occluding edges, and texture edges: A unified and robust operator based on duality," *IEEE Trans. Pattern Recognition Mach. Intell.*, vol. 20, pp. 131–146, 2008.

Callouts

AS ARCHAEOLOGICAL FRAGMENTS ARE OFTEN WEATHERED OR CHIPPED, BOTH THE ARC-

LENGTH AND THE MEASURED DIFFERENTIAL QUANTITIES OF THE CURVES ARE VERY NOISY AND WILL NOT EXACTLY AGREE.

CERAMIC POTTERY FRAGMENTS ARE COMMONLY THE MOST PROLIFIC FIND AT AN ARCHAEOLOGICAL SITE AND THEY ARE A MAJOR SOURCE OF INFORMATION CONCERNING THE FUNCTIONING OF THE SOCIETY AT THE SITE.

COMPUTATIONAL METHODS FOR OBJECT RECONSTRUCTION MUST HAVE DIGITAL REPRESENTATIONS FOR THE FRAGMENTS TO BE ASSEMBLED.

GIVEN THE IMPORTANCE OF POTTERY IN UNDERSTANDING OUR PAST, THERE HAS BEEN A SIGNIFICANT RESEARCH EFFORT PUT FORTH TO DEAL WITH THE ANALYSIS OF POTTERY FRAGMENTS.

SINCE THE DAWN OF MANKIND, OUR ANCESTORS HAVE BEEN CREATING MAN-MADE OBJECTS TO SERVE IN EVERYDAY LIFE.

[SP]