Digital Culture & Heritage
Patrimoine & Culture Numérique

Haus der Kulturen der Welt, BERLIN
Aug. 31st - Sept. 2nd, 2004
31 Août - 2 septembre 2004

DIGITAL RECONSTRUCTION OF FRAGMENTED ARTIFACTS

Helena Cristina G. Leitão and Jorge Stolfi

Fluminense Federal University, Niterói RJ, Brazil and State University of Campinas, Campinas SP, Brazil

http://www.ic.unicamp.br/~stolfi/EXPORT/projects/fragments/Welcome.html

Published with the sponsorship of the French Ministry of Culture and Communication

Actes publiés avec le soutien de la Mission de la Recherche et de la Technologie du Ministère de la Culture et de la Communication, France

Interprétation simultanée du colloque et traduction des actes réalisées avec le soutien de l'Agence Intergouvernementale de la Francophonie

Abstract (EN)

We describe an ongoing research project on efficient methods for identification and reconstruction of broken objects among large collections of irregular fragments. Applications include archaeology, art restoration, failure analysis, and other disciplines. Our solution for flat objects uses multi-scale matching and constrained dynamic programming. We are now extending it to curved pottery fragments through computational stereo vision techniques.

Keywords: Digital reconstruction, multi-scale matching, pottery fragments.

Zusammenfassung (DE)

Es wird ein noch laufendes Projekt beschrieben, bei dem es um effiziente Methoden zur Restaurierung von Objekten geht, die aus großen Ansammlungen unregelmäßiger Bruchstücke zusammengefügt werden sollen. Die Anwendungsbereiche umfassen die Archäologie, die Restaurierung von Kunstgegenständen, die Fehleranalyse, u. a. m. Unsere Lösung für flache Objekte bedient sich des Multi-Scale Matching, eine Methode zur Anpassung verschiedener Maßstäbe, und dem Constrained Dynamic Programming, der Programmierung voneinander abhängiger bewegter Teile. Wir sind dabei diese Lösung unter Einsatz der Computergestützten Stereobild-Technik auch auf gewölbte Tonwarenfragmente anzuwenden.

Schlüsselwörter: Digitale Rekonstruktion, Multi-Scale Matching, Tonwarenfragmente.

Résumé (FR)

Présentation d'un projet de recherche sur l’identification et la reconstruction d'objets brisés dans de vastes collections de fragments irréguliers. Ce probléme apparaît dans l’archéologie, la restauration d’objets d’art, l’analyse d’accidents, et bien d’autres disciplines. Notre solution pour objects plats utilise un système de couplage à échelles multiples et la programmation dynamique à contraintes. Nous cherchons maintenant à étendre le champ d'application de la méthode aux morceaux courbes de céramiques, par des techniques de vision stéréoscopique par ordinateur.

Mots clés: Reconstruction numérique, Echelles multiples, Fragments de céramiques.

I. Introduction

Plain ceramic fragments are often found by the thousands in archaeological digs. Those finds probably include enough matching pieces to build several nearly complete artifacts; however, the reconstruction is hardly attempted due to the large cost of trying millions of possible pairs by hand.

This is a problem where computers could be of great help. Here we describe a computational procedure for the reconstruction of ceramics objects from a large number of undecorated fragments. The procedure uses only the two-dimensional outline of the fragments, i.e. the shape of the fracture line. The goal is to find any pairs of curves that have long sections with similar shapes --- as similar as expected from adjacent fragments of the same object. If the two paired segments are sufficiently long, and their shapes are sufficiently similar, we can discard the hypothesis that the similarity is due to mere coincidence. This is the hard part of the problem: once the matching pairs have been found, reassembling the object --- actually or virtually --- is a fairly trivial task.

Our technique has been validated with freshly broken fragments of flat ceramics tiles. A project for testing it on curved pottery fragments from a real archaeological site is underway. The main difficulty is data acquisition, namely digitizing the fragment outlines in three dimensions with sufficient accuracy, with equipment that can be realistically used in typical museum settings. We plan to use a combination of geometric and photometric stereo to extract that information from multiple digital photographs of the fragments.

II. Related work

At present, the reconstruction of archaeological fragments is done largely by hand. Computers are used, if at all, only in the enhancement, classification, and presentation of scanned images of the fragments (Kalvin et al. 1996) which are indexed and retrieved based solely on textual descriptions provided by the user.

Computer vision and pattern matching techniques have occasionally been used to automatically extract indexing and matching information from digital images of archaeological artifacts, e.g. by Menard and Sablatnig (1997). The specific problem of identifying adjacent ceramic fragments by matching the shapes of their outlines was recently considered by Üçoluk and Toroslu (1999). Their algorithm considers only a fixed scale of resolution, and therefore has large expected asymptotic cost; no real-world tests are reported in the article. Levoy and others (1998) are investigating the use of surface matching techniques for the purpose of reassembling a fragmented marble mural. Watanabe and others (2003) have recently proposed a solution for pottery fragments, that includes surface acquisition and surface matching.

The fragment re-assembly problem is similar to that of automatic reassembly of jigsaw puzzles, which was addressed by several researchers including Burdea and Wolfson (1989). However, these techniques rely on special characteristics of puzzle pieces, such as smooth borders and sharp corners, which are lacking in archaeological materials. On one hand, in fractured objects, each fragment corner is generally incident to three fracture lines, and two of them usually make an angle close to 180º. Thus, a typical corner will be hardly detectable in one of its incident fragments. On the other hand, physical fracture lines tend to display sharp bends even where there are no ideal corners. Nevertheless, this corner-based approach has been applied recently to potsherd reconstruction by Hori, Imai, and Ogasawara (1999, 2000).

More generally, the problem can be viewed as a special case of object recognition by approximate outline matching (Pope 1994). However, in this field one typically assumes that the given outlines are to be matched against a small set of fixed templates; while, in our application, the templates are the fragments outlines themselves, which may number in the thousands. Therefore, we have to discard most standard object recognition methods because they rely on extensive preprocessing of the templates.

Multi-scale techniques have often been used for image-based and outline-based shape matching, e.g. by Witkin (1983), Mokhtarian and Makworth (1992), and Mokhtarian (1995). Most prior work in this area is based on the identification of certain critical points of the outline, such as corners or curvature extreme. As we noted above, this approach is neither feasible nor useful in the case of ceramic fragment outlines. Therefore, our algorithm applies the multi-scale approach directly to the comparison of fragment outlines, without prior identification of critical points.

III. Formal problem

The identification of matching fragments from outlines is not only an interesting archaeological problem, but a challenging *computational* one as well. The main difficulties are the large number of potential solutions, the lack of reliable breakpoints on the outlines, the non-trivial pairing of samples between independently acquired outlines, and the presence of "noise" due to material and measurement sources.

1. Potential solutions

We assume that we have N fragments, and that the outline of each fragment is digitized as a polygonal line with L sample points on the average. Instances of the problem that are of interest to archaeologists have thousands of fragments. Assuming that the fragments have not been damaged by tumbling or rough handling, details as small as 0.5 millimeter or less may be important, so for typical coarse pottery fragments L is typically on the order of 1000 or more.

When trying to fit a fragment A against another fragment B, we have to find two pieces of curve (*segments* in what follows), one on each outline, whose shapes match within a certain tolerance _ and have at least K samples; where K is proportional to log N, in practice between 10 and 100. There are about L choices for the starting point of the segment on A. For each of those choices, finding a matching segment in B is going to require a number of operations proportional to LK. Thus, the total number of operations required to all possibilities of all $N(N-1)/2$ pairs of fragments is several times N^2KL, or 10 billion even for a small collection of fragments.

2. Fragmentation

When a relatively flat friable object is shattered into many fragments, the resulting network of fractures typically has a *recursive binary partition* structure. That is, we can order the fractures in time, so that each fracture cuts across one of the pieces produced by all previous fractures, and divides that piece in two. As a consequence, at each vertex of the network, there are usually two fractures meeting with T-like geometry: the first one runs uninterrupted through the vertex, while the second one ends abruptly there. Moreover, since each fracture typically replaces an n-sided piece by two pieces having a total of $n+4$ sides, it follows that the expected number of sides per fragment is close to 4 (even though the expected number of neighbors of each fragment is close to 6, by Euler's theorem on planar maps).

3. Shape of fracture lines

In friable materials, the fracture lines are self-similar to a certain extent; that is, a small portion of a contour, magnified, is statistically similar to the whole. Indeed, fracture lines are a classical example of *fractal curves* (Falconer 1990). It is precisely their randomness at all scales that makes our multi-scale approach possible.

The fractal model is appropriate for fractures in rigid and granulated material such as ceramic, stone, and stucco. Therefore, the main applications of our work are expected to be in archaeology (pottery, clay tablets, chipped stone tools, mural paintings) and art conservation (murals, statuary). Other possible applications are paleontology, surgery, and forensics (bone fragments) and failure analysis (debris). The model may also apply to certain types of old paper or vellum, that tend to crumble rather than tear.

4. Approximate matching

The matching segments of contour are not identical, of course: they are perturbed by "noise" from several sources. During contour acquisition, some errors are introduced due to the finite resolution of the imaging device, and to effects like shadows and parallax, that act separately on the two contours. More importantly, there is often loss of material along the fracture edge. Although the fragments suffer surprisingly little wear while they are buried, they may be significantly damaged by rough handling during recovery, e.g. by sieving or by being stored in sacks or boxes without any padding.

For these reasons, comparing two segments of contour is not a trivial operation. It is not just a matter of comparing samples with some tolerance: since the noise can change the length of the curve, and therefore the number of samples, we must also allow for a flexible pairing of the samples.

5. Information contents of fracture lines

As observed above, the correct partner of a given segment of fracture line must be identified among NL other segments of outline. Therefore, on a first approximation, we need to extract about $\log_2(NL)$ bits of useful information from the segment's shape – "useful" in the sense that the same bits can be extracted from its correct partner, in spite of all sources of error.

In a previous investigation, we have found that the outlines of well-preserved ceramic fragments contain about 1.8 bits of useful shape information per millimeter (Leitão and Stolfi 2000w). Therefore, in order to find a matching partner of a fragment, we need to have about $\log_2(N\,L)/1.8$ millimeters of common boundary -- i.e., about 17 mm for $N = 1{,}000{,}000$ and $L = 1000$. Note that this minimum shared length grows only logarithmically with the collection's size. Note also that the image resolution required for this -- less than 0.5 mm per pixel -- is well within the capabilities of current semi-professional digital cameras.

IV. Multi-scale contour matching

In order to avoid comparing every part of every fragment outline against every other part --- which would be too much work even for a computer --- we use the *multi-scale paradigm* (Witkin 1983): namely, we solve several versions of the problem, at increasing scales of detail, starting with the most coarse; and at each scale we simply improve the solution found at the previous scale. At each scale we use an incremental version of dynamic programming to compare the two outlines.

1. Re-sampling at multiple scales

The first step is to prepare several copies of every fragment outline, each copy re-sampled to one half of the previous resolution. Each outline must be smoothed before re-sampling to avoid *aliasing artifacts*, which would distort its shape. We stop when the typical fracture line separating two adjacent segments has been reduced to a couple of samples. Thus, if the initial outlines have L points, on the average, we need to prepare $O(\log_2 L)$ copies of each.

2. Initial candidates

Then, among the last set of outlines, we look for points that match locally, that is, that have complementary curvatures --- one convex, the other concave, but with similar radii. Since this criterion is not very selective, it will result in a large set S of candidate matches --- perhaps thousands for each fragment. However, this step can be performed fairly quickly because each outline are reduced to a few tens of points.

3. Candidate mapping and refinement

In the next step, we locate these candidate matches on the previous set of outline copies, which have twice as many sample points. Thus each candidate becomes a sequence of 2--3 points on one outline and another 2--3 points on another outline. We compare theses two point sequences, and, if they are too dissimilar, we discard the candidate. This step is repeated several times, at increasing resolution. At each pass, we have a certain set of candidates, which are pairs of pieces of outlines that, at the current scale of detail, seem to match. We identify each candidate match on the higher-resolution outlines, we re-compare the two curves, and reject candidates which are now seen as too dissimilar.

4. Elimination criteria

The elimination criteria are adjusted so that at least half of the candidates are eliminated at each pass, and at the end we have at most 6 candidates per fragment. Since the number of sample points per candidate doubles at each step, the two factors cancel and each pass takes about the same time.

Our experiments with broken tiles show that most of the correct pairs that could be identified by an exhaustive search do get included in the initial candidate list, and survive until the end of the process. Moreover, if the surviving candidates are ranked by similarity, taking into account all the detail, the correct pairs generally end up at the top of the list.

5. Comparing two curves

As discussed in section III, when comparing two segments we need to choose a pairing of their samples. We use the pairing that makes the two segments most similar to each other, i.e. which gives the smallest total difference between paired samples, in a specific sense. There is a standard method for finding such pairings, called *dynamic programming* (DP). When comparing two segments with M samples, the basic DP algorithm takes time proportional to M^2. However, we use the multi-scale approach here, too: namely, we assume that the best sample pairing at a certain scale is not too different from the best pairing found at the previous, coarser scale. Thanks to this assumption, we are able to reduce the cost of DP from quadratic to only proportional to M.

V. Test with flat tile fragments

We have validated our program on a artificial but reasonably realistic test case. The input data for this test was a set of 112 ceramic fragments, obtained by shattering five rectangular unglazed ceramic tiles, about 25.0 cm by 6.0 cm. See figure 1.

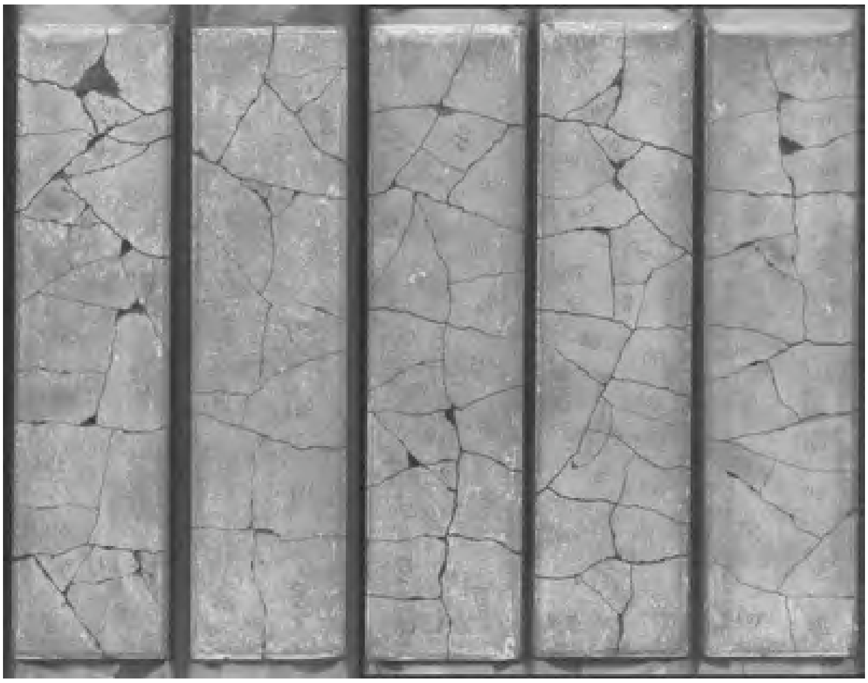

Figure 1: Test fragments, manually re-assembled.

The fragments were directly digitized with a standard flatbed scanner (UMAX model UC630 Maxcolor) at 300 pixels per inch. The flat side of each fragment was lightly rubbed with chalk to enhance contrast. The resulting images are shown in figure 2.

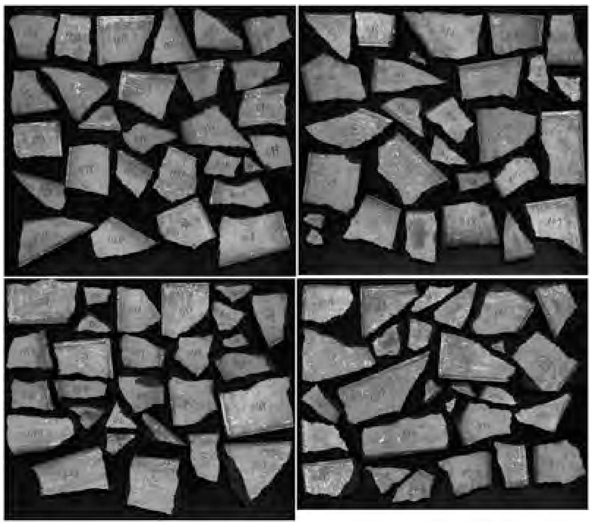

Figure 2: Input images for the test.

The outlines of these fragments were extracted with sub-pixel accuracy by a simple contour-following algorithm, and submitted to our multi-scale matching procedure. The initial set of 85509 candidates was progressively reduced through 5 steps of mapping and re-matching until only 49 pairs were left. Figure 3 shows the first 35 of these candidates, after ranking them by increasing discrepancy. Correct matches are marked by stars.

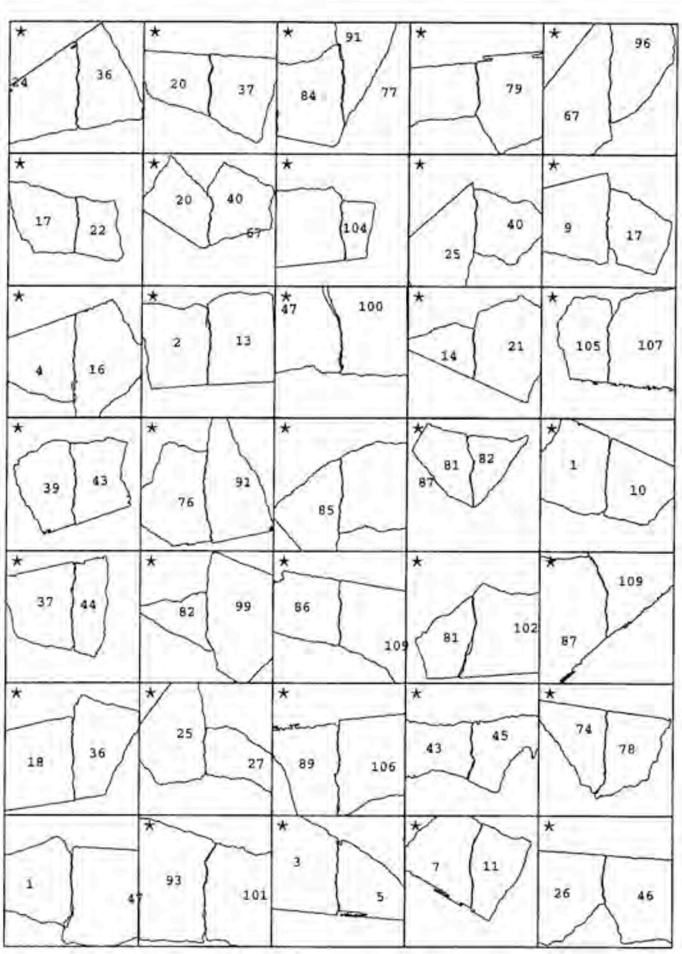

Figure 3: The 35 best pairs found by the program.

VI. Real archaeological fragments

We are currently adapting our programs to handle actual archaeological pottery fragments. As a test case, we have photographed 400 fragments recovered from a pre-Columbian site in Brazil (Caju site, in Campos, RJ) excavated by the Brazilian Archaeology Institute (IAB). The fragments come from undecorated utility pots and funerary urns, attributed to the Una ceramics tradition and dated from the 9th century CE.

1. Fragment imaging

The main difficulty is data acquisition: since the fragments are generally curved, we cannot image them with a flatbed scanner, as in our validation test. Moreover, the fracture edges of many fragments appear to have been damaged by rough handling after excavation, so we anticipate that it will be necessary to use the *medial outline* of each fragment -- a virtual profile traced along the center line of the fracture surface -- rather than the actual fracture edges.

The bottom line is that we need to obtain very precise three-dimensional models of the fragments. To that end, we took about 20 images from each fragment batch, from several viewpoints and illuminated from three different angles. We used a commercial 3-megapixel digital camera (Sony Mavica CD300) that delivers a resolution 5 pixels per mm, which we believe is accurate enough for our goals. The result was a database of about a thousand images similar to figures 4 and 5.

Figure 4.

Figure 5.

2. Image calibration and segmentation

Before we can extract the three-dimensional models from these images, they must undergo geometric and photometric calibration. For this purpose, we included in each picture a calibrated grid, as well as the reference objects and gray patches seen in figure 8.

This step is complicated by some automatic color "enhancements" made by the camera, which we weren't able to turn off completely. Then a human operator must manually select a point on each fragment image, and type the corresponding fragment number which was written in ink on each fragment, see figure 6. (Fortunately this data can be replicated automatically to all views of the same batch.) Then the individual fragment images are extracted, as in figure 7, with the *image-forest transform* (Falcão, Stolfi and Lotufo 2004).

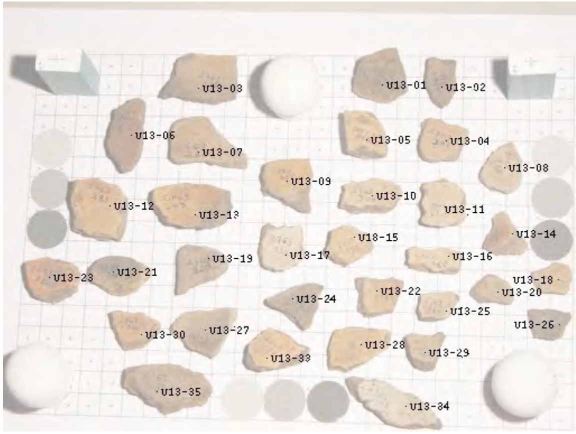

Figure 6.

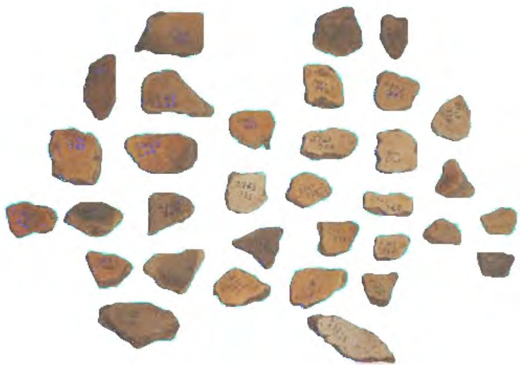

Figure 7.

3. Stereo reconstruction

To recover the third dimension with the accuracy needed for shape matching, we plan to combine two computer vision techniques, neither of which is sufficiently accurate per se. The first method is traditional *geometric stereo*, where one infers the depth of a point on the object from its position in two images taken from different angles (as in figure 4) or with the help of a mirror (as in figure 5). The accuracy of the reconstructed depth is expected to be roughly equal to the image resolution divided by the angle between the two viewing directions, namely about 0.2 mm / 0.1 rad = 2 mm for our frontal-view pairs. The second method is *photometric stereo*, where one obtains the local slope at a point of the object from the way its apparent color changes when lighted from different directions. From this information one can in principle obtain *relative* displacements between nearby points with fairly high accuracy. By combining these two pieces of data, from all 20 views of each fragment, we believe it will be possible to build a 3D model with accuracy 0.5 mm or better. Needless to say, this is a very challenging data fusion project, and we are barely starting to tackle it.

VII. Conclusions

The validation experiment with fresh fragments, although modest in size, has demonstrated the possibility of automatic of mechanically identifying adjacent fragments pairs by matching their outlines. It has also shown that the multi-scale matching procedure is an effective method for finding such pairs, since the number of surviving candidates decreases very quickly as the are re-tested with increasing resolutions. While we still do not have a program that can be used for actual archaeological puzzles, we are confident that such software will soon become available and will become a standard tool for archaeologists and museum curators.

Acknowledgments

We wish to thank Anamaria Gomide, Islene Garcia, Selma Leitão, Elaine Angrisano, Jack Sickermann, and Cid Souza for help with this manuscript. We are grateful to Dr. Paulo Seda of IAB for the opportunity to photograph the Caju fragments, and to Dr. Xavier Perrot for the invitation to this conference. This research was funded in part by the Brazilian agencies CNPq, CAPES, FAPESP, and FAPERJ.

References

Burdea G.C. & H.J. Wolfson (1989). Solving jigsaw puzzles by a robot. *IEEE- Transactions on Robot and Automation,* 5(6), 752--764.

Mokhtarian F. (1995) Silhouette-based isolated object recognition through curvature scale space. *IEEE - Transactions of Pattern Analysis and Machine Intelligence*. 17(5), 539--544.

Mokhtarian F. & Mackworth A.K. (1992). A theory of multi-scale, curvature-based shape representation for planar curves. *IEEE - Transactions of Pattern Analysis and Machine Intelligence*. 14(8), 789--805.

Falcão A. X., Lotufo R. A., and Stolfi J. (2004). The Image foresting transform: Theory, algorithms, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 26(1), 19--29.

Leitão H.C.G. & Stolfi J. (2002). A multi-scale method for the re-assembly of two-dimensional fragmented objects. *IEEE - Transactions of Pattern Analysis and Machine Intelligence*. 24(9),1239--1251.

Leitão H.C.G. & Stolfi J. (2004). Information contents of fracture lines. *International Journal of Computer Vision,* to appear.

Witkin A.P. (1983). Space-scale filtering. In *Proceedings of IJCAI''83 - 8th. International Joint Conference of Artificial Intelligence,* 1019--1021.

Wolfson H.J. (1990). On curve matching. *IEEE Transactions of Pattern Analysis and Machine Intelligence*. 12(5), 483--488.

Üçoluk G. & Toroslu I.H. (1999). Automatic reconstruction of broken 3-D surfaces objects. *Computer & Graphics*, 23(4), 573--582.

Leitão H.C.G. & Stolfi J. (2000). A multi-scale method to re-assembly of fragmented objects. In M. Mirmehdi & B. Thomas (Eds.) In *Proceeding of the British Machine Vision Conference 2000*.vol.2, 705--704.

Watanabe Y. Tanaka K., Abe N., Taki H., Kinoshita Y. and He S.(2003). Toward the automatic restoration of relics from slice image. *Proceedings of International Cultural*

Heritage Informatics Conference (*ICHIM 03*), 1--17, September 2003, Paris, France.

Menard C. & Sablatnig (1997). On finding archaeological fragment assemblies using a bottom-up design. In *Proceeding of the 21st. Workshop of the Austrian Association for Pattern Recognition*, 203--207.

Levoy M. (1999) Scanning the fragments of the Forma Urbis Romae. Consulted October, 2000.

Pope A.R. (1994). Model based object recognition: A survey of recent research. Technical Report TR-94-04, Berkeley University.

Falconer K. (1990). Fractal Geometry: Mathematical Foundations and Applications. John Willey and Sons.

Hori K., Imai M. & Ogasawara T. (2000). Hierarchical Description of a Contour for Reconstruction of Broken Earthenware. *The Trans. Electronics, Information and Comm. Engineers DII* (in Japanese), J83-D-II(5), 1392--1394.

Hori K., Imai M. & Ogasawara T. (1999). Joint Detection for Potsherds of Broken Earthenware. *Proc. IEEE Conf. Computer Vision and Pattern Recognition* (*CVPR ´99*), vol. 2, 440--446, June 1999.