

3D Pottery Shape Similarity Matching Based on Digital Signatures

Anestis Koutsoudis^{1,2} and Christodoulos Chamzas¹

¹ Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, Greece.

² Cultural and Educational Technology Institute-Research Centre 'Athena', Xanthi, Greece.

Abstract

Pottery is considered as one of the most representative categories of artifacts in the cultural heritage domain. Nowadays, several 3D digitized replicas are publicly available over the Web. The content richness of 3D pottery is of great importance in the archaeological research domain. This information can be used by special software tools that will provide the archaeologist with automated shape matching mechanisms that take under consideration the 3D morphological characteristics. In this work, we present a digital signature that can be used for pottery shape matching. The signature carries information related to the 3D morphology of a vessel and its extraction is based on a mesh pre-processing procedure. We have implemented a software tool that extracts the proposed signature and we performed initial tests on shape matching using a 3D pottery database composed of both modeled and digitized 3D models.

Keywords: 3D pottery repositories, shape matching, 3D digitization, content-based retrieval

1 INTRODUCTION

As 3D scanning technology advances, commercial systems are becoming affordable and less cumbersome. The digitization procedure is slowly transforming into a trivial task to perform. Nowadays, 3D digitization is considered a common practice within the cultural heritage domain.

Mature research domains such as computer vision and 3D computer graphics in combination with the low cost of high bandwidth Internet connections allow 3D content to be delivered easily at any low-end personal computer. Hence, 3D replicas of real world objects are becoming the new wave of multimedia content over the Web. Nevertheless, 3D replicas of artifacts are still considered a supplemental medium to digital photographs, regardless of the superior content richness they provide. It is a fact that a 3D replica is liberated from the "Do not touch!" label and allows scientists to experience a simulated physical interaction in their own time and space. Another advantage of a digital replica has to do with the ability to be studied simultaneously by several scientists. Furthermore, archaeology is a science that is based on both perception and comparisons and thus the need for identifying similar artifacts that will lead scientists to conclusions is inevitable.

The automation of identifying shape similarities between 3D replicas using a Web based search engine dictates the application of content based retrieval (CBR) mechanisms. The application of CBR in the cultural heritage domain might enable archaeologists to discover similarities and coherences between typologies in a more efficient way. It might also allow the reduction of the time required to come to conclusions. Moreover, it is considered as an approach to overcome the multi-language textual annotation

barriers experienced when using a keyword-based search engine.

CBR is a very active research field and numerous papers propose methodologies for identifying shape similarities for both two dimensional data (e.g. digital images) as well as three dimensional data. Several methods introduce the idea of encoding global or local shape feature distributions into histograms,¹ skeletal graphs,² spherical harmonics,³ and geometrical moments.⁴

We considered ceramics as a good application for our research not only because it is one of the most representative categories of artifacts exhibited in museums but also because ceramics enjoy a great interest from both scholars and general public. They have a remarkable continuity through time, allowing archaeologists to understand the society that produces them.

To this end, we propose a method for 3D pottery shape matching. The signature extraction is performed after

¹Robert Osada et al., "Matching 3D Models with Shape Distributions," paper presented at the International Conference on Shape Modeling and Applications, Genova, Italy, May 7–11, 2001.

²Alexander Agathos et al., "Retrieval of 3D Articulated Objects Using a Graph-based Representation," paper presented at the Eurographics 2009 Workshop on 3D Object Retrieval, Munich, Germany, March 29, 2009.

³Michael Kazhdan and Thomas Funkhouser, "Harmonic 3D Shape Matching," paper presented at the SIGGRAPH 2002 Technical Sketches, San Antonio, TX, USA, July 2002.

⁴Athanasios Mademlis et al., "3D Content-based Search-based on 3D Krawtchouk Moments," paper presented at the Third International Symposium on 3D Data Processing, Chapel Hill, NC, USA, June 14–16, 2006.

the completion of a two step pre-processing phase. During this phase a 3D vessel is normalized in terms of scaling and orientation (pose). The algorithm relies solely on one of the most important characteristics of vessels, their axial symmetry. The morphological features of the vessel are extracted using multiple plane-based contouring along the vessel's axis of symmetry.

The subsequent sections of this paper are organized as follows. In section two we describe the proposed scale and pose normalization phases, while in section three we describe the pre-processing and the feature extraction phases. Then we continue with some initial performance results that were performed on a 3D database repository.¹ Finally, we conclude by giving some thoughts on the future development of the current shape matching system.

2 MESH PREPROCESSING PHASE

The 3D mesh preprocessing is performed into two distinct steps. These are the scale and pose normalization phases. Figure 1 illustrates a 3D replica of an ancient Greek lekythos before and after the two normalization phases. Each phase is described in detail in sections 2.1 and 2.2.




Figure 1. Scale and pose normalization.

2.1 SCALE NORMALIZATION

Initially, the 3D model is positioned in an arbitrary position within the 3D space. Once the scale normalization phase is completed, the vessel is limited within a unit bounding sphere. The scale normalization of the vessel is based on the computation of the minimum bounding sphere as described by Emo Welzl.² Welzl uses the basic idea of Seidel's linear programming algorithm to compute the minimum bounding sphere of a point set in 3D space in linear time. The minimum bounding sphere is computed in an incremental way, starting with an empty set and then adding points one after the other while maintaining the smallest enclosing sphere for the points considered at

¹Anestis Koutsoudis et al., "A 3D Pottery Database for Benchmarking Content Based Retrieval Mechanisms," paper presented at the Eurographics 2008 Workshop on 3D Object Retrieval, Crete, Greece, April 15, 2008.

²Emo Welzl, "Smallest Enclosing Disks," *Lecture Notes in Computer Science* 555 (1991): 359.

each time. The center coordinates and the radius of the bounding sphere are calculated using as an input the mesh vertices coordinates. Then the appropriate spatial affine transformations are applied so that the minimum bounding sphere has a radius equal to a unit bounding sphere (0.5 units) and its center lies on the origin of a left-handed 3D Cartesian coordinate system. It should be mentioned that although the scale normalization phase restricts the vessel within the limits of the unit bounding sphere, its axis of symmetry (V_a) is still arbitrarily oriented.

2.2 POSE NORMALIZATION

A common characteristic of all vessels, even for those that are not wheel-made, is their rotational body. The pose normalization phase exhibits this property to detect the axis of symmetry (V_a). Several approaches have been proposed in the past in order to estimate the axis of symmetry.³ Our method's goal is to rotate the vessel within the bounding sphere so that V_a is set parallel to the Y axis of the 3D Cartesian coordinate system (C_y) and the top of the vessel is oriented towards the positive side of C_y . In the usual case, where the vessel carries a handle, another rotation around the C_y is also performed so that the handle is positioned towards the positive side of the C_x axis. In cases where the vessel carries more than one handle, the one most distant to V_a is identified and oriented again towards the positive side of the C_x axis. Furthermore, real vessels are far from perfect and thus they are considered as *noisy surfaces of revolution*. This is due to the many imperfections caused during the production phase and by the subsequent erosion. Thus, the algorithm attempts to converge to an optimum V_a and then set it parallel with C_y .

In order to perform the previously described affine transforms, the 3D mesh is segmented by using multiple, vertical to C_y , plane-based contouring and then discriminating at each contouring level between the parts that belong to the vessel's main body (MB group) and those that belong to handles or feet (HF group). The objects that belong to the MB group are further divided into those that belong to the outer shell of the main body and those that belong to the inner shell. Additionally, each contour can carry from none up to several objects, depending on the morphology of the vessel. If, for ex-ample, a single object is found in all

³Yan Cao et al., "Geometric Structure Estimation of Axially Symmetric Pots from Small Fragments," paper presented at the International Conference on Signal Processing, Pattern Recognition and Applications, Crete, Greece, June 25–28, 2002; Chaouki Maiza et al., "SemanticArchaeo: A Symbolic Approach of Pottery Classification," paper presented at the Seventh International Symposium on Virtual Reality, Nicosia, Cyprus, October 30–November 4, 2006; Hubert Mara et al., "Orientation of Fragments of Rotationally Symmetrical 3D-shapes for Archaeological Documentation," paper presented at the Third International Symposium on 3D Data Processing, Chapel Hill, USA, June 14–16, 2006.

contouring levels, this indicates that the vessel carries no handles or feet.

More specifically, the initialization of the pose normalization phase is done by performing *Principal Components Analysis* (PCA) using the mesh vertices and by rotating the vessel so that the first principal axis P_a is set parallel to C_y . Nevertheless, it is known that PCA is highly affected by the mesh density and the vertices distribution. For example, in situations where a 3D laser scanner has been used, the 3D mesh is composed by variable vertex density areas depending on the complexity of the surface. This is a result of the mesh optimization and data reduction required for increasing the real time visualization frame rate. Non-uniform vertex distribution can also be found in vessels that have been 3D reconstructed using the popular *shape-from-silhouette* method. In some cases, a solution to this problem is to insert vertices on the mesh. The number of additional vertices is determined by the area covered by each triangle or polygon. Even then, PCA results are still not able to be trusted as a stand-alone approach to identify V_a correctly. Additionally, PCA results are also affected by vessels which carry a single handle or a set of highly asymmetric handles. Despite the previous limitations, PCA is used in our approach as an initialization step that reduces the total number of recursions performed by our algorithm to achieve the correct orientation of the vessel. Figure 2 depicts cases of erroneous calculation of V_a by using only PCA.




Figure 2. Performing Principal Component Analysis.

Once PCA is completed, the algorithm rotates the vessel so that $P_a \parallel C_y$. However, for some vessel shapes, P_a is not the axis with the closest direction to the V_a axis. Figure 2 depicts such a case. Thus there is a need to identify which one of the three principal axes, P_a , P_b or P_c , is the one that is closest to the direction defined by the V_a axis. In order for the algorithm to perform this task, the axial symmetry of the vessel is exploited. Multiple plane-based contouring is performed along all three principal axes (P_a , P_b , P_c). Then *circular regression* is computed for all objects. The circle fitting error (variance) of each object is summed up into three totals (one for each axis). The lowest variance is expected to be found along the principal axis that is closer to the V_a axis. This is due to the circular characteristics of vessels and their axial symmetry property. A vertical planar contour on an axis which is closer to V_a will produce objects with a significantly lower fitting error, as their shape will be

in most cases imperfect or incomplete circles. Hence, if the lowest variance is not found on P_a , then the appropriate rotation is performed so that the principal axis (P_b or P_c) with the lowest variance is set parallel to C_y .




Figure 3. Contouring along each principal axis.

Figure 3 illustrates the objects that appear after contouring the mesh along each one of the three principal axes. It is obvious that in the first case (P_a) the objects appearing at each contouring level are imperfect circles, while in the two other cases (P_b , P_c) they are not.

The algorithm then initiates a recursive phase, which involves mesh contouring, object identification, temporary axis of symmetry detection, and vessel rotation. This recursion will converge on detecting the optimum axis of symmetry (V_a) and meet its main goal ($V_a \parallel C_y$). After each mesh contouring and in order to avoid the noisy transitional areas on the vessel's surface (e.g. regions where a handle begins to extend from the main body), we consider a contouring level as useful if and only if the previous and the next to the current contouring level have the same number of objects. Those that fulfill the previous condition are considered as useful for detecting V_a .

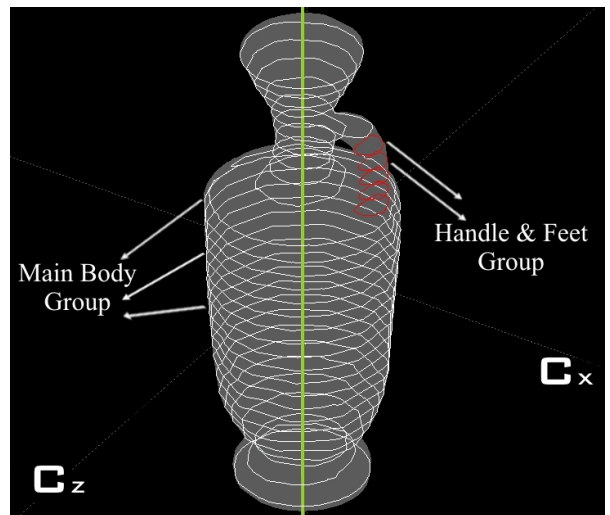


Figure 4. Object grouping.

The objects that appear at each contouring level are identified and characterized as parts of the MB (inner and outer shell subclasses) or the HF group (fig. 4). The MB group subclass segmentation is compulsory.

Although most of the 3D scanners cannot reach the inner bottom side of a vessel, they are still able to capture the inner surface around the rim; thus the discrimination between the inner and outer shell is mandatory. The outer shell objects of the MB group of each useful contouring level are used for estimating the axis of symmetry. Circular regression is computed only for the objects that belong to that subclass. The center coordinates of the fitted circles are used by the *Singular Value Decomposition* (SVD) factorization method in order to determine the equations of a line in the 3D space. This line is considered as the vessel's *temporary axis of symmetry* (V_{at}) (fig. 5). The inclinations of the 2D projections of V_{at} against two 2D planes C_x - C_y and C_z - C_y are calculated and used to rotate the vessel accordingly around the C_z and C_x axes so that $V_{at} \parallel C_y$. The algorithm repeats the previous steps until the resulting rotations are below a given threshold (in our implementation it is set to 0.01°) or after reaching a maximum number of iterations, which indicates the algorithm's oscillation. Finally, we consider that $V_{at} = V_a$, and thus we accept that the best convergence to get $V_a \parallel C_y$ has been achieved.

The next step of the pose normalization phase is performed only when the vessel is identified as carrying handles or feet. If such is the case, the vessel is rotated around the C_y axis at a given angle. The angle is calculated between the C_x axis line segment and a line segment that starts from V_a and ends at the center of gravity of the object most distant to V_a located at any contouring level. This orientation correction completes the optimum pose setting where apart from setting $V_a \parallel C_y$, the handle is now located at an *optimum* position to be picked up by a right-handed person. In cases where two opposite handles exist, the result is that both handles are positioned along the C_x while in other special cases where multiple handles are present the most distant one is considered as the dominant one; its center of gravity will be considered as the end of the previously described line segment. Figure 6 illustrates examples of such cases.

Finally, in order to ensure that the top of the vessel is oriented towards the positive side of C_y , two orthographic depth map images are taken from two virtual cameras with opposite view points (fig. 7). The virtual cameras are positioned on the positive and negative sides of the C_y axis, outside of the unit bounding sphere limits, allowing a parallel projection of the 3D world on a 2D square plane (with side size equal to 1 unit). The depth map image that contains pixels with grayscale values closest to zero (colors closer to black define greater distances between the virtual camera and the object's surface) in a circular area around the axis of symmetry (V_a) is probably the one which was taken by the virtual camera that aims towards the inner side of the vessel through its rim. Although this approach has been proven to perform quite well through our tests, there are cases of failure, especially when the 3D model is produced using the

shape-from-silhouette approach, which is known for its inability to reconstruct concave areas that are not visible on the object's silhouette curve. If the top of the vessel is identified on the negative side of C_y , a 180° rotation around the C_x is performed.




Figure 5. Detecting the axis of symmetry.

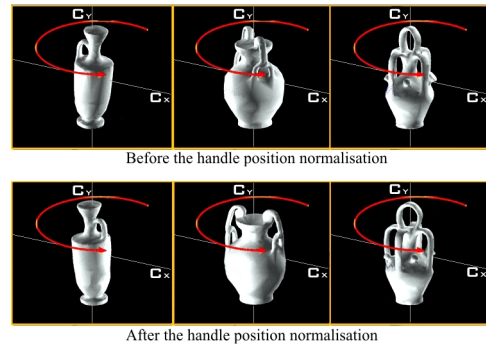


Figure 6. Handle position normalization.




Figure 7. Identifying the vessel's top.

3 SIGNATURE EXTRACTION

After the completion of the scale and pose normalization, the final set of sequential contours along

the C_y are used to extract the signature. The objects that appear at each contouring level are once again identified as parts of the MB group or parts of the HF group. Circular regression is again performed for every object in all contouring levels. Thus, properties such as the center of gravity of the fitted circle, its radius, and the circular fitting error are calculated. The signature (descriptor) we propose in this work is divided into two different parts. The first part is used to encode the body of the vessel while the second encodes the handles and the feet. This approach has been selected in order to be able to perform partial shape matching between vessels using only their body or handles and to allow the usage of different weights when matching the body or the handles.

The first part of the descriptor is the *Profile Vector*, a 1D vector which carries the radiuses of the best fitted circle of the objects that belong to the outer shell of the vessel's main body. A graph depiction of the vector is visually similar to a quantized version of the vessel's body outer profile.

The second part of the descriptor encodes the properties of all objects that belong to HF group in a 2D array. This time, the position of the best fitted circle of each object is quantized into a predefined area-sector around the unit circle. In our tests we selected $\pi/8$ as an angle increment step that results in 16 discrete sectors (fig. 8).

Additionally, the radius of each object and its distance from V_a are also held in the same array in a data structure. In our experiments the dimensions of this array are 32 rows by 16 columns. Each row represents a contouring level, while each column indicates the position of the HF object around the unit circle at a contouring level.

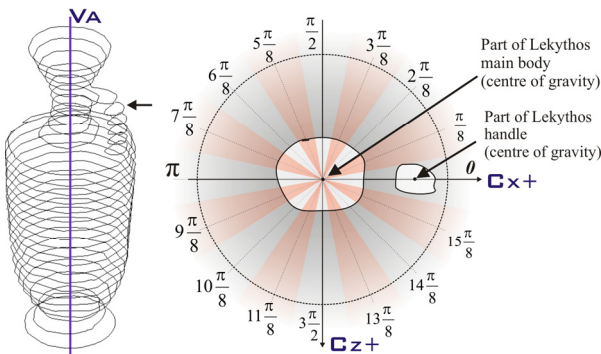


Figure 8. Quantizing the position of an object's gravity center to one of the predefined sectors.

4 ALGORITHM IMPLEMENTATION AND INITIAL SHAPE MATCHING TESTING

We developed a stand-alone application in Borland Delphi. The software can visualize in real time the pose

normalization phase. An optimized implementation of the minimum bounding sphere algorithm, written in C++ capable of handling large numbers of points has been used.¹

The total number of sequential contouring levels used in our experiments was 32. This number was selected because it results in a dense distribution of contouring levels within the unit bounding sphere that are able to capture the detailed features of our testbed database models. The descriptor extraction time has been reduced by introducing an initial mesh decimation step. This is an important step for digitized models, as their large amounts of data introduce delays to the contouring procedure and the application of the affine transforms.

We tested our approach on a 3D pottery repository with 1,276 models (94 digitized and 1,182 manually 3D modeled). They cover several categories, such as ancient Greek, Native American, ancient Roman, modern pottery and abstract symmetrical objects. A total of 1012 vessels were successfully oriented, resulting in an 85.9% success rate. The current implementation fails with shapes like very shallow plates, due to the low number of contouring levels. The average run time for the pose normalization and descriptor extraction is 25.6 seconds on an AMD Athlon running at 2.2 GHz with the real time algorithm operation visualization enabled.

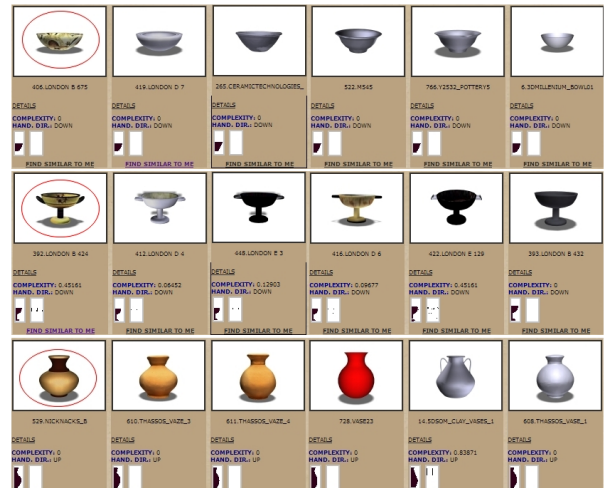


Figure 9. Query-by-example results.

For our initial testing, the extracted features were encoded into an XML schema and organized into a native XML database. Several retrieval tests were performed comparing only the profile vectors by using the Euclidean distance as a similarity metric. Figure 9 illustrates promising results of several queries-by-example. In each row the first 3D model is the query object. At the moment, we are calibrating our

¹Bernd Gartner, "Fast and Robust Enclosing Balls," *Lecture Notes in Computer Science* 1643 (1999): 325.

repository in order to compute detailed precision-recall performance histograms for several types of shapes including ancient Greek vessels.

5 CONCLUSIONS

In this paper, we presented our initial work on a novel descriptor for 3D pottery content-based retrieval, which is based on a scale and pose normalization algorithm. We implemented the algorithm as a stand-

alone application and attempted some initial tests on shape matching that look promising. At the moment, we are working on a handle matching metric which will introduce the idea of fuzziness (matching between neighboring sectors on the unit circle) and the descriptor encoding as a grayscale image in order to apply image based shape matching metrics. We are also working on vessel segmentation by constructing convex hulls derived by the MB and HF objects.

ACKNOWLEDGEMENTS

This paper is part of the 03ED679 and 03ED698 research projects, implemented within the framework of the “Reinforcement Programme of Human Research Manpower” (PENED) and co financed by National and Community Funds (25% from the Greek Ministry of Development-General Secretariat of Research and Technology, and 75% from the EU-European Social Fund). The authors would like to acknowledge and thank the Cultural Heritage Unit of the Cultural and Educational Technology Institute/R.C. ‘Athena’ for their support in this work, Prof. Robert Jacob from the Computer Science Department of Tufts University, and the Centre of Advanced Spatial Technologies of the Hampson Museum (<http://hampsonmuseum.cast.uark.edu>) for providing us with their 3D vessel collections.

BIBLIOGRAPHY

- Agathos, Alexander, et al. “Retrieval of 3D Articulated Objects Using a Graph-based Representation,” paper presented at the Eurographics 2009 Workshop on 3D Object Retrieval, Munich, Germany, March 29, 2009.
- Cao, Yan, et al. “Geometric Structure Estimation of Axially Symmetric Pots from Small Fragments,” paper presented at the International Conference on Signal Processing, Pattern Recognition and Applications, Crete, Greece, June 25–28, 2002.
- Gartner, Bernd. “Fast and robust enclosing balls.” *Lecture Notes in Computer Science* 1643 (1999): 325–338.
- Kazhdan, Michael, and Thomas Funkhouser. “Harmonic 3D Shape Matching,” paper presented at the SIGGRAPH 2002 Technical Sketches, San Antonio, Texas, USA, July 2002.
- Koutsoudis, Anestis, et al. “A 3D Pottery Database for Benchmarking Content Based Retrieval Mechanisms,” paper presented at the Eurographics 2008 Workshop on 3D Object Retrieval, Crete, Greece, April 15, 2008.
- Mademlis, Athanasios, et al. “3D Content-based Search based on 3D Krawtchouk Moments,” paper presented at the Third International Symposium on 3D Data Processing, Chapel Hill, USA, June 14–16, 2006.
- Maiza, Chaouki, et al. “SemanticArchaeo: A Symbolic Approach of Pottery Classification,” paper presented at the Seventh International Symposium on Virtual Reality, Nicosia, Cyprus, October 30–November 4, 2006.
- Mara, Hubert, et al. “Orientation of Fragments of Rotationally Symmetrical 3D-shapes for Archaeological Documentation,” paper presented at the Third International Symposium on 3D Data Processing, Chapel Hill, USA, June 14–16, 2006.
- Osada, Robert, et al. “Matching 3D Models with Shape Distributions,” paper presented at the International Conference on Shape Modeling and Applications, Genova, Italy, May 7–11, 2001.
- Welzl, Emo. “Smallest Enclosing Disks.” *Lecture Notes in Computer Science* 555 (1991): 359–370.