

## 3D Pottery Shape Similarity Matching Based on Digital Signatures

Anestis Koutsoudis, Christodoulos Chamzas  
akoutsou@ceti.gr, chamzas@ceti.gr

Department of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, 67100, Greece

### Abstract

Pottery is considered as one of the most representative category of artifacts in the cultural heritage domain. Nowadays, several 3D digitized pottery replicas are available to the public through the Web. The content richness of 3D pottery replicas are 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 which will take under consideration the three dimensional morphological characteristics provided by 3D replicas. In this work, we present a digital signature (descriptor) that can be used for shape matching and carries information regarding the 3D morphology of a vessel. The signature extraction is based on a mesh pre-processing procedure. We tested our approach on a 3D pottery database composed by 3D models both modeled and 3D scanned with multiple methodologies such as laser scanning triangulation, shape from stereo and shape from silhouette.

### 1. Introduction

3D replicas of real world objects are becoming the new wave of multimedia content over the Web. As 3D scanning technology advances and commercial systems are becoming cheaper and less cumbersome, the digitisation procedure is transforming into an easier task to perform. This evolution has also affected the application of 3D scanning in the cultural heritage domain by becoming slowly a common practice. Additionally, research domains such as computer vision and computer graphics, which are considered as mature research fields, in combination with the low cost of high bandwidth Internet connections allow 3D content to be delivered easily at any average personal computer. Although 3D models are still considered as a supplemental medium to digital photographs, the content richness they provide equip archaeologists with a high detail 3D replica which is liberated by the “*Do not touch*” label and at the same time can be studied by multiple scientists at their own time by simulating a physical interaction with the real artifact using real time 3D graphics visualisation techniques.

Furthermore, archaeology is a science that is based on both perception and comparisons and thus the need for identifying similar artifacts is inevitable. In order to automate the procedure of identifying similarities between 3D replicas the need of applying content based retrieval (CBR) mechanisms is more than obvious. It might enable archaeologists to discover more efficiently similarities and coherences between typologies by reducing the total time required to come to conclusions or to overcome the multi-language textual annotation barrier. We considered ceramics as a good applicant for our research not only because it is one of the most representative categories of artifacts exhibited in museums but also because they enjoy a great interest from both scholars and general public, they have a remarkable continuity through time allowing archaeologists to understand the society that produce them and also because they are considered as “*3D scanner friendly*” digitisation candidates.

As shape matching is an important element for CBR, recognition and classification, we propose a method for the extraction of a digital signature which can be used in 3D pottery shape matching. The signature extraction is performed after the completion of two preprocessing normalisation phases related to the scale and the pose of a vessel in the 3D space. The morphological features of the vessel are extracted

using multiple plane-based contouring along the vessel's axis of symmetry. The axis of symmetry is detected solely based on one of the most important characteristics of pottery, its axially symmetric shape.

The subsequent sections of this paper are organized as follow: In sections two we describe the proposed scale and pose normalisation 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][2]. Finally, we conclude by giving some thoughts on the future development of the current shape matching system.

## 2. Preprocessing of the vessel's 3D mesh

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

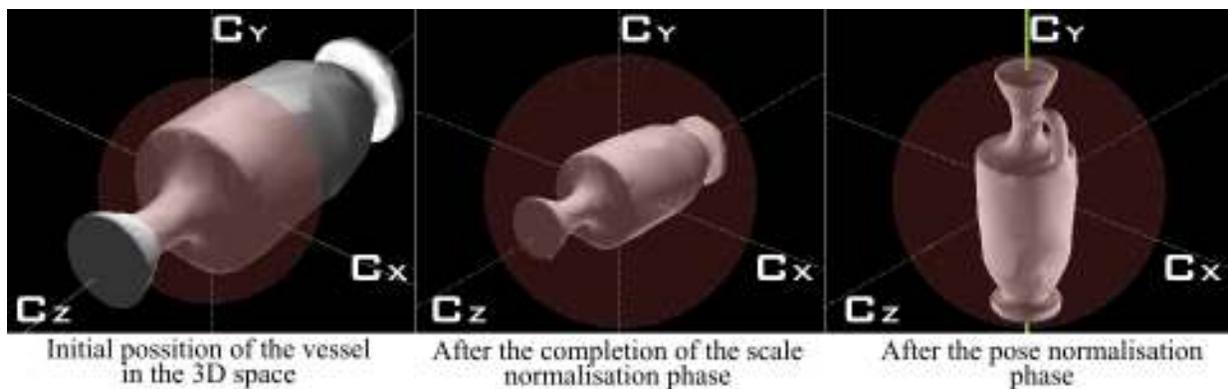


Figure 1: Pose normalisation of a lykethos 3D replica.

### 2.1 Scale Normalisation

Initially, the 3D model is positioned in an arbitrary way within the 3D space. Once the scale normalisation phase is completed, the vessel is limited within a unit bounding sphere. The scale normalisation of the vessel is based on the computation of the minimum bounding sphere as described by Emo Welzl's [3]. 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 each time [3]. Therefore, the properties of the bounding sphere such as its centre coordinates and radius are calculated using as input the coordinates of the mesh vertices. Then, the appropriate spatial transformations are performed so that the minimum bounding sphere has a radius equal to a unit bounding sphere ( $0,5$  units) and its centre lies on the origin of a 3D Cartesian coordinate system. It should be mentioned that although, the scale normalisation phase restricts the vessel within the limits of the unit bounding sphere, its axis of symmetry ( $V_a$ ) can still be arbitrary oriented.

### 2.2 Pose Normalisation

A common characteristic for all vessels, even those which are not wheel made, is that carry a rotational body. The pose normalisation phase exhibits the property of axial symmetry. Several approaches have been proposed in the past in order to estimate the axis of symmetry [4]-[6]. Our method is based on a

multiple level mesh contouring approach. Its goal is to rotate the vessel within the bounding sphere so that its  $Va$  is set parallel to the  $Y$  axis of the left handed 3D Cartesian coordinate system ( $Cy$ ) and the top of the vessel is oriented towards the positive side of  $Cy$ . In the general case where the vessel carries a handle then another rotation around the  $Cy$  is also performed so that the handle is oriented towards the positive side of the  $Cx$  axis. In cases where the vessel carries more than one handle or feet then the most distant to the axis of symmetry is identified and oriented again towards the positive side of the  $Cx$  axis. Furthermore, real vessels are far from perfect and thus they are considered as noisy surfaces of revolution. This is due to many imperfections caused during the production phase and also by the subsequent erosion. Thus, the algorithm attempts to converge to an optimum  $Va$  and then set it parallel with  $Cy$ . In order to perform the previously described affine transforms, the 3D mesh is segmented by using multiple, vertical to  $Cy$ , plane contours and then discriminate in each contour 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 mode divided to those that belong to the outer shell of the main body and the inner shell. Additionally, each contour can carry from none up to several objects depending on the morphology of the vessel. If for example, a single object is found in all contours, this indicates that the vessel carries no handles or feet.

More specifically, the initialization of the pose normalisation phase is done by performing *Principle Components Analysis (PCA)* using the vertices of the mesh and by rotating the vessel so that the first principal axis  $Pa$  is parallel to  $Cy$ . Nevertheless, it is known that *PCA* is highly affected by the vertices density and their distribution on the surface of the 3D mesh. In situations where a 3D laser scanner has been used for the acquisition of the vessel's mesh, we usually experience areas with variable vertices densities depending on the complexity of the surface. The purpose behind this is the mesh optimization which will result reduction on the data being handled by the graphics card and thus increase the visualisation frame rate. Non-uniform vertex distribution can also be found in vessels that have been digitized using the popular *shape-from-silhouette* method. In some cases, a solution to this problem is given by adding vertices on the surfaces defined by each triangle. The number of vertices is determined by the area covered by each triangle. Even in such cases, the result of the *PCA* is still inadequate to be used as a stand alone metric to identify the  $Va$  correctly. Additionally, *PCA* is proved to be inadequate in cases where the vessel carries a single handle or highly asymmetric handles. Despite the previous problems, in our approach *PCA* is only used as an initialization step that actually reduces the total number of recursions required to achieve the correct pose of the vessel. The following figure depicts cases of erroneous calculation of the  $Va$  by relying only on *PCA*.

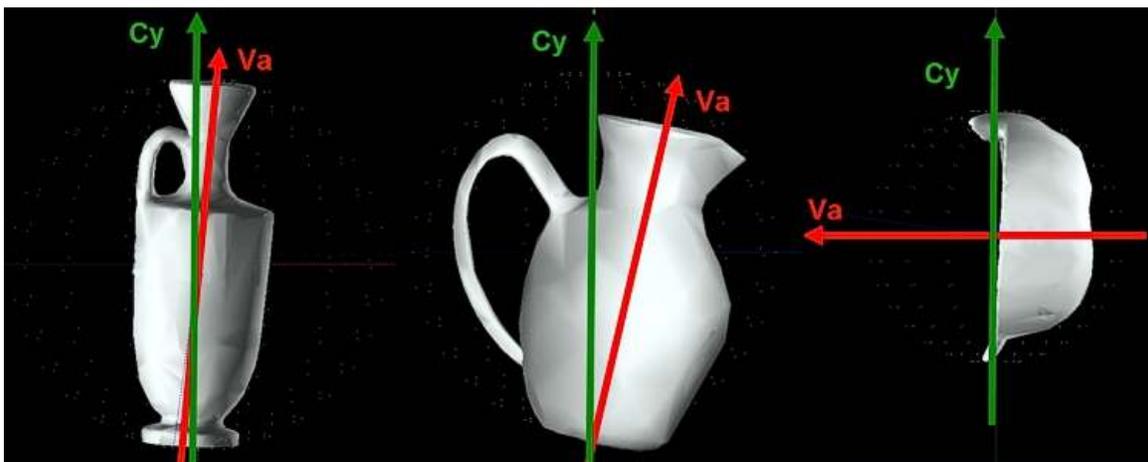


Figure 2: Performing Principal Component Analysis

Once *PCA* is completed, the algorithm rotates the vessel so that  $Pa \parallel Cy$ . However, for some vessel shapes,  $Pa$  is not the axis with the closest to the direction of  $Va$ . Figure 2 depicts such a case. Thus, there is a need to identify which one of the three principal axes  $Pa$ ,  $Pb$  or  $Pc$  is the one that is closest to the direction defined by  $Va$ . In order for the algorithm to perform this task, the axial symmetry of the vessel is once again exploited. Multiple plane based contouring is performed along all three principal axes ( $Pa$ ,  $Pb$ ,  $Pc$ ). In each case, the planes are vertical to one of the axes. The objects that appear at a contour are identified and characterized as parts of the *MB* (inner and outer shell subclasses) or the *HF* group (Figure 3). Then circular regression is computed only for those that belong to the outer shell subclass of the *MB* group. The *circle fitting error* (variance) that is resulted by each object is summed up for each principal axis into three totals. The lowest variance is expected to be found along the principal axis which is closer to  $Va$ . 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  $Va$  will result in objects with less circular regression variance as their shape will be in most cases imperfect or incomplete circles. Hence, if the lowest variance is not found on  $Pa$ , then the appropriate rotation is performed so that the principle axis with the lowest variance is parallel to  $Cy$ .

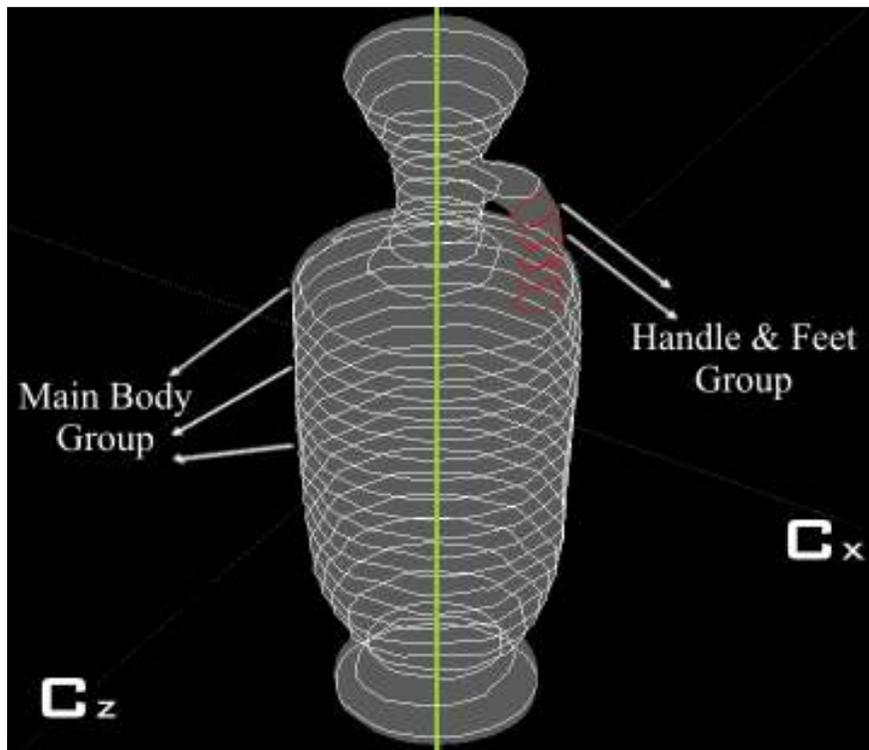


Figure 3: Performing object grouping for each contour

Figure 4 illustrates the objects that appear after contouring the mesh along each of the three principal axes. It is obvious that in the first case ( $Pa$ ) the objects that appear at each contouring level are imperfect circles while in the two other cases ( $Pb$ ,  $Pc$ ) are not. In cases where  $Pb$  or  $Pc$  is found to be the correct axis, an appropriate rotation is performed to get  $Pb$  or  $Pc$  parallel to  $Cy$ .

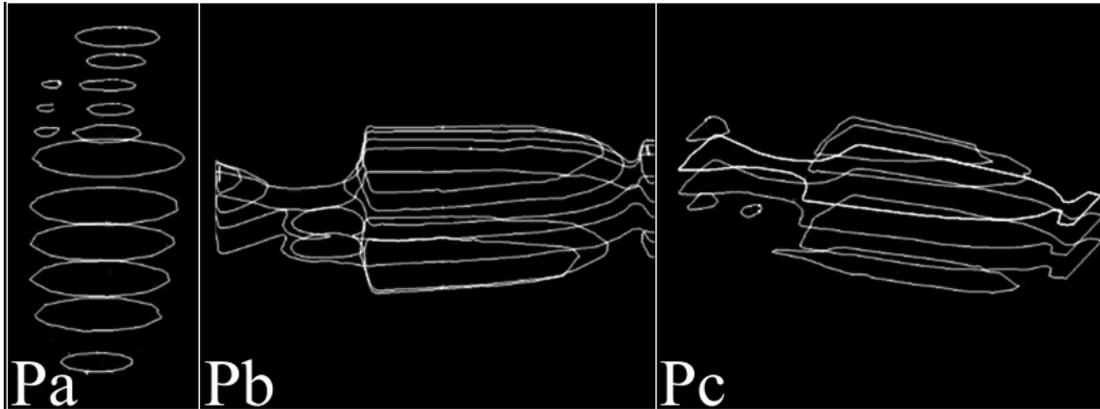


Figure 4: Detecting the axis with the closest to  $V_a$  direction

The algorithm then continues into 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 contouring of the mesh 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 contour as *useful* if and only if the previous and the next to the current contour have the same number of objects. Those that fulfill the previous condition are considered as useful for detecting  $V_a$ . The outer shell objects of the MB group of each useful contour are used. Circular regression is computed for the object belonging to that subclass. The centers of the fitted circles are used with 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* ( $V_{at}$ ) (Figure 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 resulted rotations are below a given threshold (In our implementation it is set to  $0,01^\circ$ ) or after reaching a maximum number of iterations which indicates algorithm oscillation. Finally, we consider that  $V_{at} = V_a$ , and thus we accept that the best convergence to  $V_a \parallel C_y$  has been achieved.

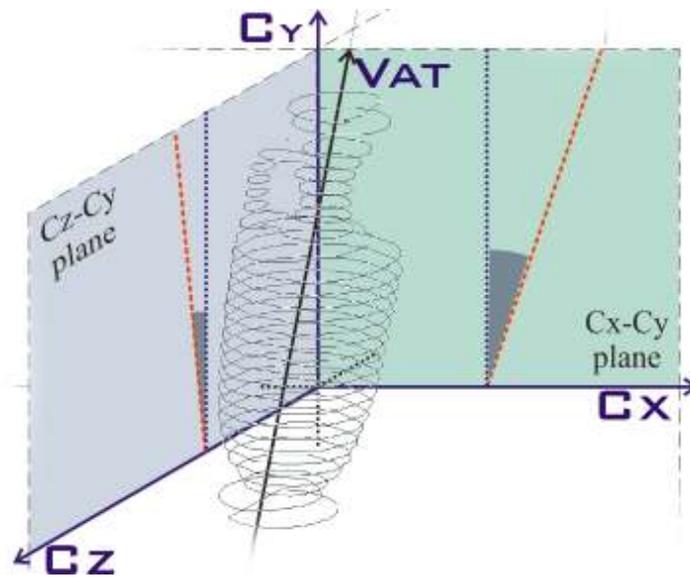


Figure 5: Estimating  $V_{at}$  using Singular Value Decomposition

The next step of the pose normalisation phase is performed only when the vessel is identified to carry handles or feet. If such is the case, the vessel is rotated around the  $C_y$  axis at a given angle which is equal to the angle between the  $C_x$  axis and a line segment that starts from  $Va$  and ends at the centre of gravity of the most distant to  $Va$  object of any contour. This pose correction completes the optimum pose setting where apart from that  $Va \parallel C_y$  is already true, the handle is located at the *optimum* position to be picked up from a right handed person. In cases of two opposite handles, the result is that both handles are positioned along the  $C_x$  while in other special cases where multiple handles exist the most distant one is considered as the dominant and its centre of gravity will be considered as the end of the line segment for the previous angle calculation. Figure 6 illustrates examples of such cases.

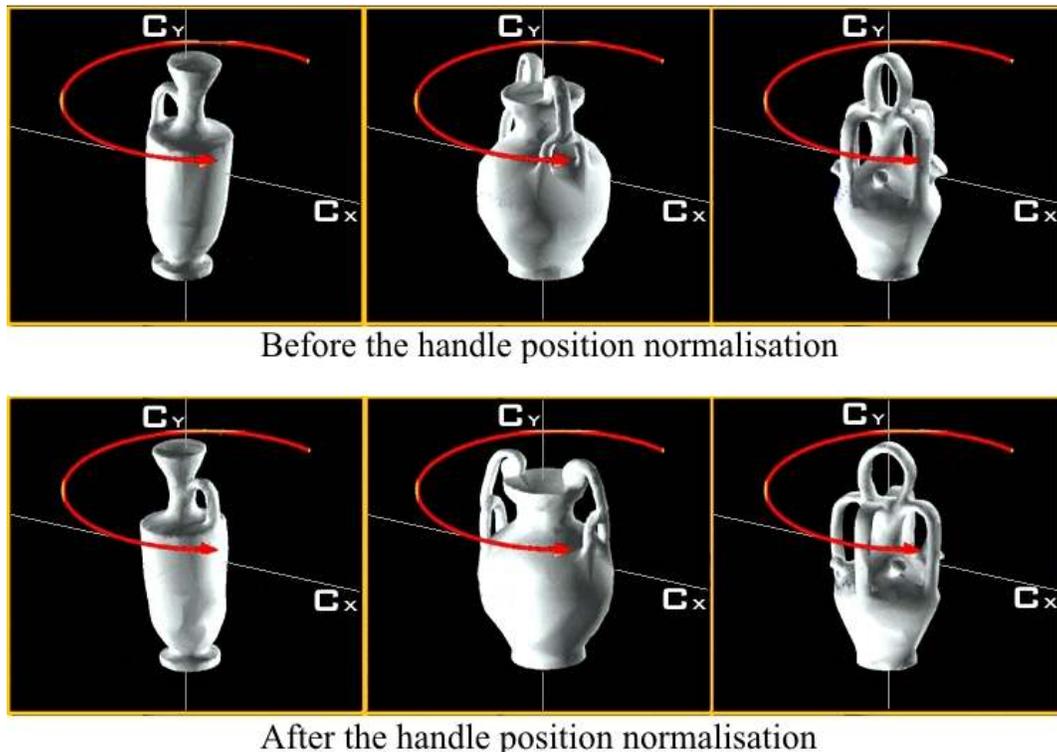


Figure 6: Different cases of handle position normalisation

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 (Figure 7). They 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 greyscale values closest to zero (Colours closer to black define higher distance between the virtual camera and the object's surface) 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 by the *shape-from-silhouette* methodology which is known for its inability to reconstruct concave areas not visible on the silhouette's curvatures. If the top of the vessel is identified on the negative side of  $C_y$ , a  $180^\circ$  rotation around the  $C_x$  is performed.

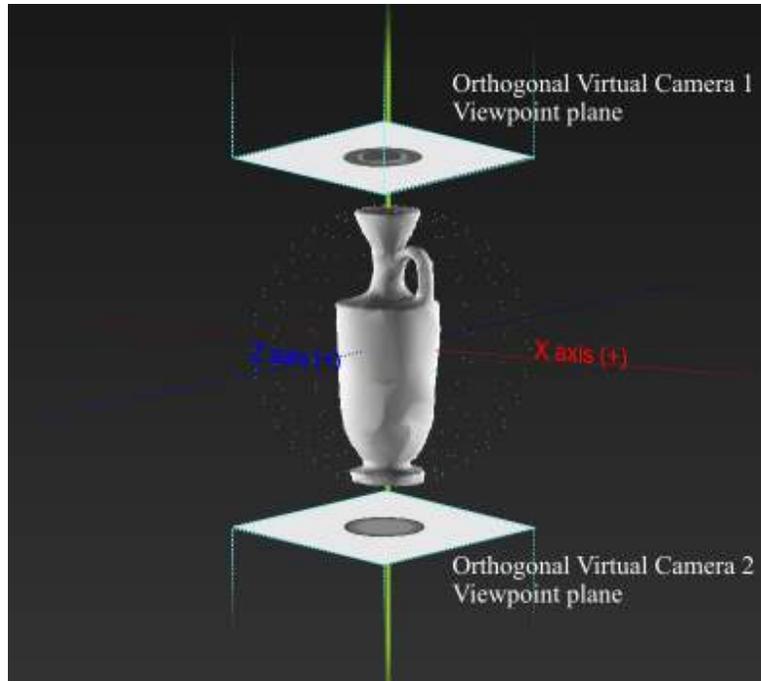


Figure 7: Identifying the top of the vessel using orthographic depth map projections

### 3. Digital Signature Extraction

After the completion of the scale and pose normalisation, a final set of sequential contours along the  $C_y$  are computed. The objects that appear in each contour 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 centre of gravity of the fitted circle, its radius and the fitting error are calculated. The  $MB$  objects are further more divided to those belonging to the outer and inner shell of the mesh. Although most of the 3D scanners cannot reach the inner bottom side of a vessel they still can capture the inner surface around the rim, thus the discrimination between the inner and outer shell is also mandatory.

The signature (descriptor) we propose is divided into two different parts. The first is used to encode the body of the vessel while the second encodes the handles and the feet. The first part, called the *Profile Vector*, is a 1D vector which carries the radiuses of the best fitted circle on the objects that belong to the outer shell of the vessel's main body at each contouring level. 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 step that resulted 16 discrete sectors (Figure 8). Additionally, the radius of each object and its distance from  $V_a$  are also hold in the same array in a data structure.

Furthermore, two metrics are proposed to be used as query filters in order to improve the retrieval performance. The *Vessel Complexity Metric (VSM)* is a ratio between the total number of  $HF$  objects and the total number of contour planes. A vessel with  $VSM=0$  indicates a vessel with no handles, thus vessels with high  $VSM$  can be filtered out. A calibration of overlapping  $VSM$  value areas is mandatory in order for the filter to improve performance while keeping the recall rate high. On the other hand, a 3D vector starting from the origin of the coordinate system and finishing at the vessel's centre of gravity can be used

as an indicator towards the direction of the position of where the handles on the vessel are located. This is true because the centre of gravity tends to move closer to regions with higher vertices density. This filter can be used for further sorting of the retrieved objects.

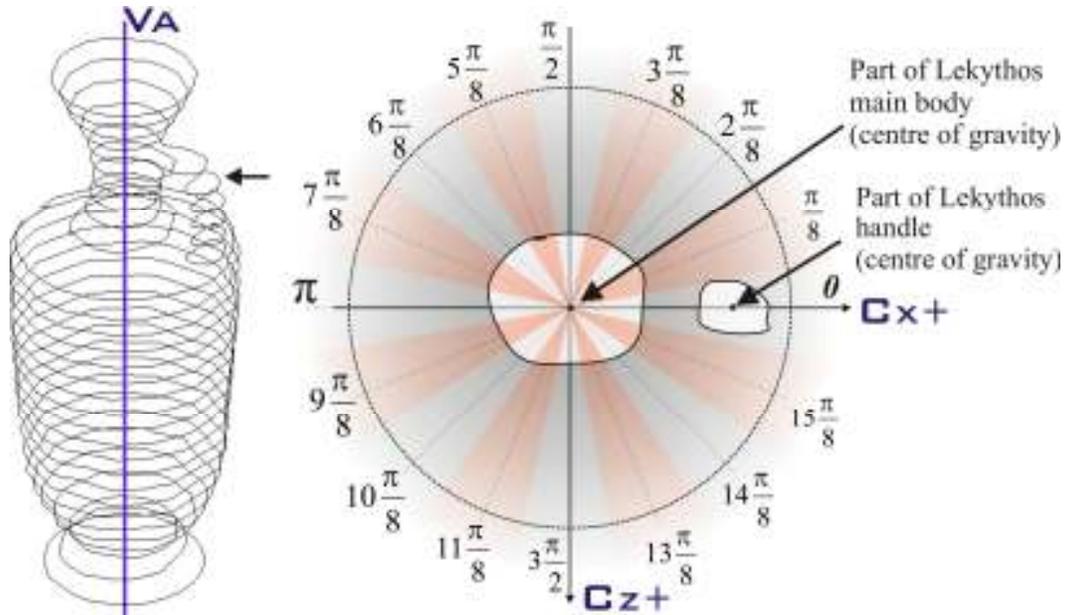


Figure 8: *Quantizing the position of an object's gravity centre to one of the predefined sectors.*

#### 4. Algorithm implementation and initial shape matching testing

We developed an application with real time visualisation of the algorithm operation in Borland Delphi using the *GLScene* and *VTK* libraries. An optimized implementation of the minimum bounding sphere algorithm, written in C++ capable to handle large number of points has been used [7]. The total number of sequential contours used in our experiments was 32 and it was selected as it results a dense distribution of contours capable to capture the features of our testbed database models. In order to improve the descriptor extraction time we introduced the idea of decimating the models. We tested our approach on a 3D pottery repository with 1,276 models (both 3D digitized and manually 3D modeled). They cover several categories such as Ancient Greek, Native American, Ancient Roman, modern pottery and abstract symmetrical objects. A total of 1,012 vessels were successfully pose normalized giving an 85,9% success to the current implementation. The current implementation fails with shapes like very shallow plates due to the low number of plane contours. The average run-time for the pose normalisation and descriptor extraction is 25,6 seconds on AMD Athlon at 2,2Ghz with the algorithm operation visualisation enabled.

For our initial testing, the extracted features were encoded into an XML schema. Several retrieval tests were performed comparing the profile vectors using Cumulative Euclidean distance as a similarity metric. Figure 9 illustrates promising results of several example based queries. In each case the first 3D model is the query object. At the moment, we are calibrating our repository in order to compute detailed precision-recall performance histograms for several types of shapes including ancient Greek vessels.



Figure 9: Query-by-example results

## 5. Conclusions

In this work, we presented our initial work on a novel descriptor for 3D pottery content based retrieval which is based on a pose normalisation method. We implemented the algorithm as a stand alone application and attempted some initial tests 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 encoding of the descriptor on a greyscale image in order to apply image based matching metrics. We are also working on vessel segmentation by using convex hulls derived by the *MB* and *HF* objects and extracting higher level semantics for each vessel.

## Acknowledgements

This paper is part of the 03ED679 research project, 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 E.U.- 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 for sharing their 3D vessel collection.

## References

1. Anestis Koutsoudis et al, "A 3D Pottery Database for Benchmarking Content Based Retrieval Mechanisms", (paper presented at the workshop on 3D object retrieval of the Eurographics 2008 annual conference, Crete, Greece, April 15, 2008).
2. Anestis Koutsoudis et al, "A 3D Pottery Content Based Retrieval Method", (paper presented at the workshop on 3D object retrieval of the Eurographics 2009 annual conference, Munich, Germany, March 29, 2009).
3. Emo Welzl, "Smallest Enclosing Disks," *Lecture Notes in Computer Science* 555 (1991):359-370.
4. Yan 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).
5. 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).
6. 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).
7. Bernd Gartner, "Fast and robust enclosing balls," *Lecture Notes in Computer Science* 1643 (1999).