




Available online at

 www.sciencedirect.com

Elsevier Masson France

 www.em-consulte.com



Original Article

3D Pottery content-based retrieval based on pose normalisation and segmentation

Anestis Koutsoudis^{a,b,*}, George Pavlidis^b, Vassiliki Liami^b,
 Despoina Tsiafakis^b, Christodoulos Chamzas^a

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

^b Cultural and Educational Technology Institute/Research Centre 'Athena', Tsimiski 58, Xanthi 67100, Greece

ARTICLE INFO

Article history:

Received 5 June 2009

Accepted 2 February 2010

Available online 19 March 2010

Keywords:

Content-based retrieval

Pose normalisation

Shape matching

3D pottery

PHP

XML

MPEG-7

Ancient Greek pottery

Pottery shape variation

ABSTRACT

This paper presents a novel compact shape descriptor designed specifically for content-based retrieval of complete or nearly complete 3D vessel replicas. The descriptor consists of two vectors that carry morphological features of the vessel's main body and appendages. The extraction of the descriptor is based on a pose normalisation preprocessing phase which is designed for axially symmetric objects. In order to evaluate the efficiency of the descriptor, we created a calibrated ground-truth database of 1012 3D digitised and manually modelled vessels and performed multiple query-by-example experiments. We present the performance of our descriptor in relation to the performance of the MPEG-7 3D shape spectrum descriptor. Additionally, a web-based 3D content-based retrieval prototype system has been developed based on open source technologies.

© 2010 Elsevier Masson SAS. All rights reserved.

1. Introduction

Pottery is one of the most representative categories of archaeological artefacts. They are exhibited in museums worldwide, enjoying a great interest from both scholars and the general public [1–3]. The large number of pottery findings in combination with the cognitive and behavioural information they carry exerts a large influence on understanding the society that produced them [4]. Modern archaeological research projects related to pottery are characterised by the substantial utilisation of digital media. Apart from excavation notes, photos and physicochemical analytical techniques, archaeologists utilise the data provided by 3D digitisation to enhance their research potential [5]. The growing interest in the 3D digitisation of pottery for preservation and dissemination purposes contributes to the generation of new digital repositories that require sophisticated indexing and management mechanisms [6]. The features and content richness of 3D replicas have to be taken into account when 3D content-based retrieval mechanisms (3DCBR) are to be developed. Replacing keywords

with a 3D mesh introduces to the user an intuitive way to depict his/her constraints in mind and at the same time allows the use of the data as the query medium [7]. 3DCBR might enable archaeologists to improve existing typologies, facilitate the identification of unknown shapes or even to discover supraregional coherencies by overcoming the multilanguage text barriers and missing standards [8]. Cultural heritage is a significant domain where explicit 3DCBR descriptors should be developed in order to provide solutions to specific classification problems [7,9].

The work presented in this paper aims at this direction. It concerns the application of 3DCBR mechanisms in the cultural heritage domain, with the main contribution being the development of a novel shape descriptor designed specifically for complete or nearly complete 3D pottery. The proposed shape descriptor encodes the morphological attributes of the vessel's main body and appendages such as handles, feet and lugs into two separate feature vectors allowing partial shape matching. Another contribution of this work is the pose normalisation algorithm on which the shape descriptor extraction relies. The algorithm consists of the scale and the pose normalisation phases which are used to create a common frame of reference for all vessels. A stand alone software tool has been implemented to extract the descriptor from 3D triangular mesh data while offering an editor with XML encoding features for annotating vessels with additional archaeological metadata. In addition, a calibrated ground-truth repository with a total of 1012

* Corresponding author.

E-mail addresses: akoutsou@ceti.gr (A. Koutsoudis), gpavlid@ceti.gr (G. Pavlidis), vliami@ceti.gr (V. Liami), tsiafaki@ceti.gr (D. Tsiafakis), chamzas@ceti.gr (C. Chamzas).

3D vessels (94 digitised and 918 modelled) has been created. The data produced have been organized into an open source native-XML database and an experimental web-based 3DCBR system using PHP technology has been developed. We have evaluated the performance of the system using our descriptor and presented it in relation to the MPEG-7 3D shape spectrum (SSD) descriptor.

The subsequent sections of this paper are organized as follows: in Section 2, we outline several related projects that have attempted to facilitate the work of archaeologists. Then, we continue with a detailed description of the preprocessing and descriptor extraction phases. In Section 4, we present our calibrated ground-truth 3D pottery repository and we continue in Section 5 by discussing several attributes of the descriptor extraction implementation. In Section 6, we evaluate the performance of our descriptor and we conclude in Section 7 by commenting on future developments.

2. Related work

Archaeologists characterise pottery using typological classification (similarities in morphological features) and attribute analysis (manufacturing, materials, treatment, etc) [25]. Although the two techniques are considered complementary, the morphological features are among the fundamental properties by which vessels are characterised [4]. At the moment, archaeologists do not have a standard software tool to assist their research and thus they follow the well-established approach of visiting museums and libraries in person [7,8,10].

Attempts have been made by researchers to provide archaeologists with tools that can either automate or semi-automate the classification problem by proposing shape matching techniques applied on 2D profiles or 3D digital replicas. Beazley's web accessed archive is considered by archaeologists as one of the most important information repositories which carries textual information, photographs and drawings of pottery, gems and sculptures [11]. On the other hand, Maiza and Gaildrat [12] proposed a symbolic approach applicable to pottery classification which has been tested on a specific vessel shape category (*sigillèles*) that carry no appendages. Maiza and Gaildrat [13] described an algorithm applicable for shape matching of both 3D fragments and complete vessels. They also proposed the use of high-level descriptions for vessel indexing. Liu De-zhi et al. [14,15] presented the idea of fitting nonuniform rational B-spline curves to 3D vessel replicas in order to produce a shape descriptor. For their experiments, they have used a group of digitised Native American vessels with no appendages. Gilboa et al. [4] proposed a methodology for typological analysis of vessel profiles and cross-sections based on curvature functions.

Furthermore, researchers have also focused on the important problem of reassembling an object by matching arbitrary fragments. Huang et al. [16] designed a system for automated reassembly of broken 3D digitised fragments that belong to a specific artefact. Sablatnig et al. [17,18] proposed an algorithm for the alignment of rotationally symmetric 3D surfaces without using corresponding points. Their algorithm can be used in a fragment matching and pottery classification system. Melero et al. [19] pro-

posed an efficient 3D vessel reconstruction and drawing system that accepts digitised 3D fragments as an input. Cao and Mumford [20] have also worked on the problem of reconstructing a 3D vessel using the curvature of a 3D fragment to detect its axis of symmetry.

Research has also been performed on the development of visualisation methodologies. Shiao et al. [5] presented a solution to the information visualisation problem of *focus-plus-context* which allows the user to focus on an individual 3D vessel replica without losing the overall context of the collection to which it belongs.

3. 3D pottery descriptor extraction analysis

In this section, we describe in detail the phases of *3D mesh preprocessing* (scale and pose normalisation), *feature extraction* and *descriptor encoding*. During the 3D mesh preprocessing, a vessel is scaled in a unit bounding sphere and a coarse detection of the vessel's axis of symmetry is performed using principal components analysis. Using mesh contouring, the vessel is segmented to the parts that belong to its main body and its appendages. Then, by using only the main body parts, the detected axis of symmetry is refined and set parallel to the Y-axis of a right-handed 3D Cartesian coordinate system. Subsequently, the vessel is rotated so that its top and appendages are placed on predefined positions. The normalised mesh is contoured again and feature extraction is initiated by performing circular regression on the objects that appear in each contouring level. Finally, features such as the radius and the centre coordinates of the best-fit circles are encoded into the proposed compact shape descriptor.

3.1. 3D mesh preprocessing

The 3D mesh preprocessing is required in order for the extracted morphological features to be comparable. It consists of the scale (SNP) and the pose (PNP) normalisation phases. Fig. 1 illustrates a 3D replica of an Ancient Greek Lekythos arbitrarily positioned and oriented in 3D space, followed by two interim states which were the result of the affine transformations of SNP and PNP.

Initially, the 3D vessel is found arbitrarily positioned within a left-handed 3D Cartesian coordinate system with its axis of symmetry (V_a) also arbitrarily oriented (Fig. 1a). The SNP is used to scale and translate the vessel so that it fits perfectly within the limits defined by a unit bounding sphere. In order to perform this task (Fig. 1b), the properties of the current minimum bounding sphere of the object are computed according to Welzl's approach [27]. Then, the appropriate scaling and translation transformations are applied to the model. All further rotations of the vessel are performed around the origin of the coordinate system which now coincides with the centre of the unit bounding sphere.

After the scale normalisation phase, PNP is initiated. Its main goal is to identify the vessel's axis of symmetry (V_a) and set it parallel to the Y-axis of the Cartesian coordinate system (C_y). As real vessels carry imperfections that have appeared during their production or by subsequent erosion, the proposed algorithm attempts to converge to an optimum V_a . Additionally, in the general case




Fig. 1. Scale and pose normalisation of an Ancient Greek Lekythos.

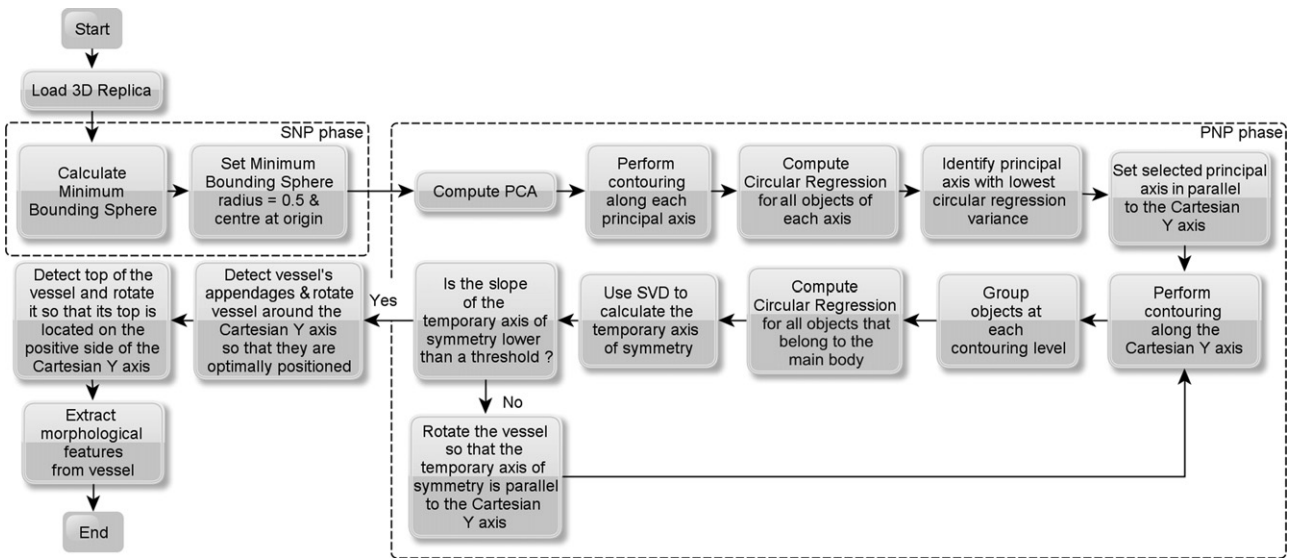


Fig. 2. Scale normalisation (SNP) and pose normalisation (PNP) flowchart.

where a vessel carries a handle, a rotation around C_y is performed so that the handle is positioned in the positive side of C_x (Fig. 1c). In cases where the vessel carries more than one appendage parts, the most distant to V_a is considered as the dominant one and it is oriented accordingly on C_x . Furthermore, the top of the vessel is detected and it is oriented in the positive side of C_y (Fig. 1c).

Fig. 2 depicts a detailed flowchart of the SNP and PNP phases.

More specifically, PNP is initiated by performing Principal components analysis (PCA) using the vertices' coordinates of the 3D mesh. It is known that PCA is affected by the appendages of a vessel and the nonuniform surface distribution of vertices, which is a common situation in digitised models [28]. Thus, it cannot be trusted as a robust approach to identify V_a . Nevertheless, it is used as an initialisation step of the proposed algorithm which results a coarse detection of the vessel's axis of symmetry. Fig. 3 illustrates three cases where the first principal axis (P_a) varies from V_a .

There are also cases where P_a is not the closest to V_a . Fig. 3c indicates such a situation. Once PCA is completed, the algorithm identifies which of the three principal axes (P_a, P_b, P_c) is closer to V_a . For performing this task, the axial symmetry of the vessel is exploited. Multiple level plane-based contouring is performed vertically along all three principal axes. Then, circular regression [29] is computed for all the objects (2D closed or open curves that match the vessel's surface) that appear at each contouring level. The circle fitting errors (variance) of each object are summed up into three totals, one for each principal axis. The lowest mean variance is expected to be found along the principal axis which has a direction close to V_a . It is known that a plane contour, which is vertical to an axis with a direction close to V_a , results in objects with low circle fitting errors as in most cases their shape is imperfect or incom-

plete ellipses. Fig. 4 illustrates an example of contouring along each of the three principal axes. In the case of Lekythos, the lowest variance is detected in P_a as the three totals are $P_{a-error} = 0.02783$, $P_{b-error} = 0.4698$ and $P_{c-error} = 0.4411$. The algorithm performs the appropriate rotation so that the principal axis with the lowest variance is set parallel to C_y .

Then the algorithm initiates a recursive phase which involves the following steps:

- multiple level plane-based mesh contouring along C_y ;
- object grouping (segmentation) at each contouring level;
- temporary axis of symmetry (V_{at}) detection;
- application of the appropriate rotations so that $V_{at} \parallel C_y$.

After each mesh contouring session, the algorithm discriminates between the objects that belong to the vessel's main body group and those that belong to the appendages group. The objects that belong to the main body group can be identified as C_y passes through them. Additionally, they are divided into those that belong to the outer-shell of the vessel's main body and those which belong to the inner-shell. This discrimination is mandatory even in the case of a closed vessel where, although a 3D scanner cannot capture its interior completely, it can still reach the inner side of the rim. Fig. 5 illustrates an example of the object grouping approach applied at a contouring level near the rim of an Ancient Greek Amphora.

The outer-shell objects are used to identify the temporary axis of symmetry (V_{at}). 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 the outer-shell objects of a contouring level important for detecting V_a if and only if the previous and the

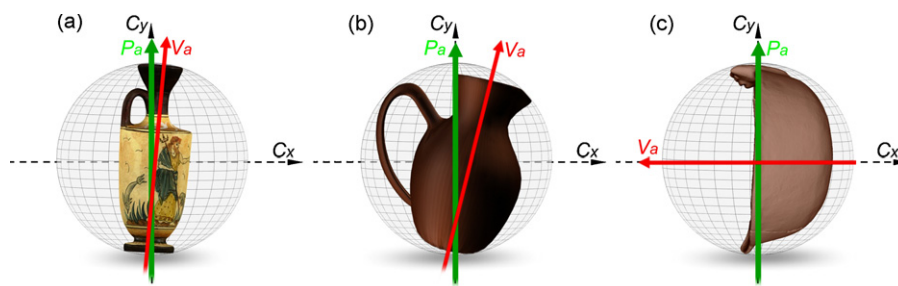


Fig. 3. V_a and P_a axes of different vessel shapes.




Fig. 4. Detecting the principal axis with the lowest mean circular regression variance.

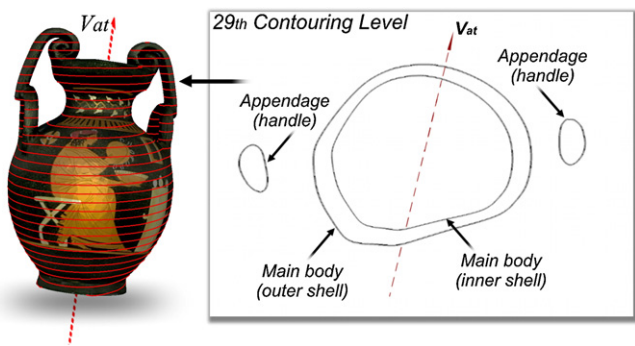


Fig. 5. Grouping objects at a contouring level.

The algorithm repeats all four previously described steps until the slopes of V_{at} are below a given threshold (In our implementation, it was set at 0.01°) or after reaching a maximum number of iterations which actually indicates a situation of algorithm oscillation. The latter appears in cases of digitised vessels where the axis of symmetry cannot be strictly defined. At the end of the process, the last V_{at} is considered as V_a .

The next step of PNP is carried out only if the vessel is identified to carry appendages such as handles, feet or lugs. In such cases, the vessel is rotated around the C_y -axis at an angle φ , such that the centre of the most distant from V_a object is placed on the C_x axis (Fig. 7). The distance of an object is equal to the distance from point O (lies on V_a) to point C (centre of best-fit circle) plus the radius of the best-fit circle (point M). Thus, the total distance is given by $d = |OC| + |CM|$.

In cases where two opposite appendages exist (e.g. two handles), the result is that both of them are positioned along the C_x axis, while in other cases (e.g. multiple handles), the most distant object is considered as dominant. Fig. 8 illustrates such cases.

The correct top-bottom vessel orientation that follows is based on the fact that the top of a vessel is expected to be hollow. Hence, in order to ensure that the top of the vessel is oriented towards the positive side of V_a , two depth map images are captured using two virtual orthographic cameras (Fig. 9), positioned on the positive and negative sides of V_a , outside of the unit sphere. These cameras allow the orthographic projection of the 3D world onto two 2D square planes which are perpendicular to V_a . The region of interest in each depth map is a predefined circular area around V_a (a circu-

next to the current contouring level contain the same number of objects. Circular regression is computed only for the selected main body outer-shell objects. Then, Singular value decomposition (SVD) is applied to the centres of the best-fit circles in order to determine the equation of a line in 3D space that connects all circle centres. This line is the current V_{at} (Fig. 6). The slopes of the 2D orthogonal projection of V_{at} onto the C_x - C_y and C_z - C_y planes are then calculated. These slopes are used to rotate the vessel accordingly around C_z and C_x so that $V_{at} \parallel C_y$.




Fig. 6. Estimating the temporary axis of symmetry (V_{at}) slopes in relation to planes C_x - C_y and C_z - C_y .

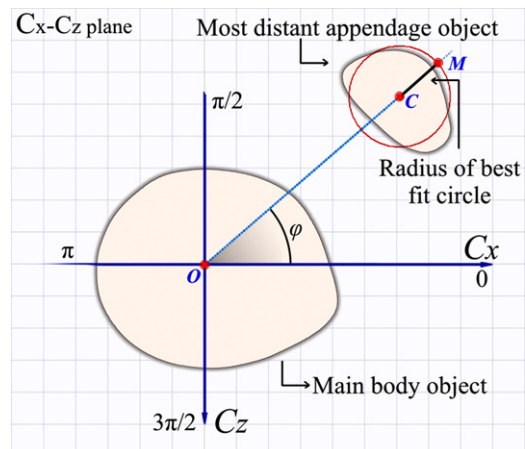


Fig. 7. Detecting angle φ of the most distant appendage object.

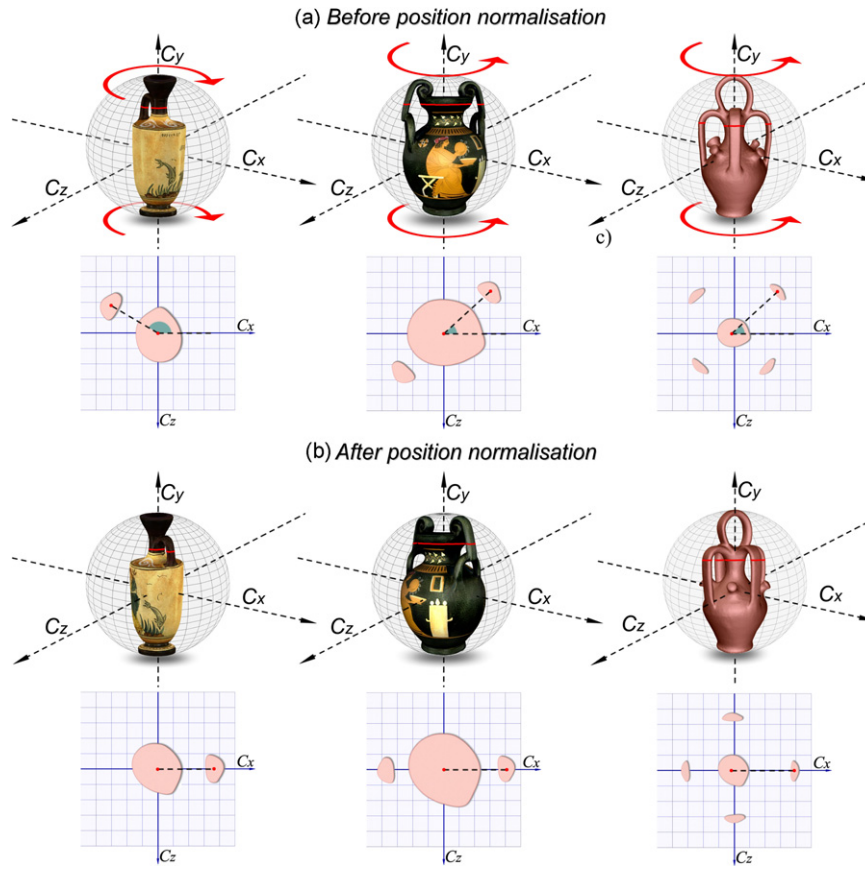


Fig. 8. Different cases of appendage position normalisation.

lar area with radius 0.1 units was used in our implementation). In the regions of interest, colours closer to black represent a greater distance between the virtual camera and the object's surface thus the region with values closer to zero indicate the camera that aims towards the inner side of the vessel. If the top of the vessel is identified by the camera on the negative side of V_n , a 180° rotation around C_x of the vessel is carried out.

3.2. Feature extraction and descriptor encoding

After the completion of the SNP and PNP, feature extraction is performed. The features are extracted from the objects produced

by the last mesh contouring step of PNP's recursive phase. These objects are grouped and circular regression is performed for every object in all contouring levels. Thus, properties such as the coordinates of the centre of the best-fit circle, its radius and the fitting errors are known for all objects. The algorithm discards the objects that are composed of two vertices (e.g. line segments produced by contouring orphaned triangles) or those which their best-fit circle has a radius smaller than a given threshold (set at 0.09 units in our experiments) as they are considered noise.

These extracted features are the basis of the proposed *Vessel Main Body and Appendages* descriptor (VMBA). VMBA is divided into two parts: the *Profile feature vector* (PFV) and the *Appendages feature vector* (AFV). PFV is a 1D vector which carries the radiuses of the best-fit circles of all objects that belong to the outer-shell of the main body group. The length of this vector is equal to the number of contouring levels and it is a quantized version of the vessel's body outer profile (Fig. 10a). On the other hand, AFV is a 2D array which encodes the positions of the appendages in a binary form. The number of rows of the 2D array is equal to the number of contouring levels while the number of columns defines the sectors around the unit circle at each contouring level (Fig. 10b). In our tests, we selected $\pi/8$ to be the increment angle that defines 16 sectors on the unit circle. The centre coordinates of each best-fit circle that belongs to the appendages group are quantized to the closest sector. Hence, if an object belongs to a given sector at a given contouring level then that AFV cell value is set to one.

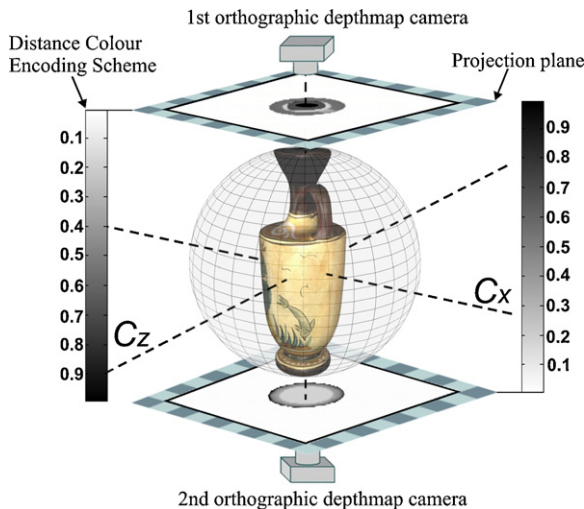


Fig. 9. Detecting the top of a vessel using orthogonal projection depth map images.

4. A calibrated ground-truth 3D pottery repository

The evaluation of content-based retrieval mechanisms is performed using ground-truth databases. Currently, there is no such database, with cultural heritage digitised 3D content, publicly avail-

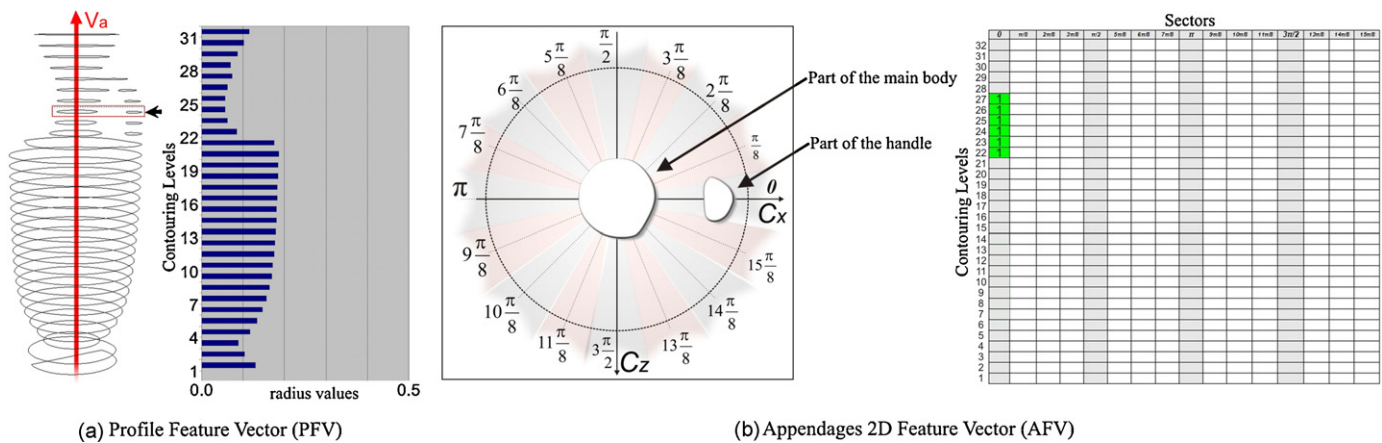


Fig. 10. Components of the vessel main body and appendages (VMBA) descriptor.

able. Thus, in order to evaluate our shape descriptor, we had to create a calibrated ground-truth dataset of 3D vessels. The 3D data are a combination of royalty free digitised and manually modelled 3D vessels, copyrighted digitised vessels provided by other research groups, manually modelled vessels based on photographs and vessels created by our 3D vessel random generator [30]. The dataset covers several shape categories such as Ancient Greek (Alabastron, Amphora, Hydria, Kantharos, Lekythos, Psykter, etc), Native American (Jar, Effigy, Bowl, Bottle, etc), modern pottery and others. Those that fall into the last two categories are considered as noise and are used to enhance the objectivity of the evaluation. The total number of vessels is 1012 (94 digitised, 718 manually modelled by different research groups and 200 generated by our random generator [30]).

The dataset was calibrated by annotating each vessel with metadata that cover both archaeological and computer graphics aspects. The archaeological metadata contain attributes such as the shape and the type of shape, the preservation state, the manufacturing technique that has been used, its attribution, the workshop's name, etc. On the other hand, the computer graphics metadata cover aspects such as the 3D digitisation technique that has been used and the availability of texture map information.

We selected five Ancient Greek shape categories for performance evaluation because of their relatively large population in the database. We have also computed their intraclass shape variation in order to support the validity of the performance results. A virtual camera positioned on C_z was used to capture a front-view depth map image for all scale and pose normalised vessels. The use of depth map images allows the calculation of the shape variation by considering not only the outline of the vessel but also its 3D surface. The corresponding pixel values of all images that belong to a given shape category were summed up and produced images such as those found in Table 1. The intraclass shape variation is calculated using the following formula:

$$\text{Shape}_{\text{variance}} = \frac{\text{Total number of nonbackground pixels with value } \leq \text{threshold}}{\text{Total number of nonbackground pixels}}$$

A pixel with zero value is considered as background. The threshold (set at 0.8 in our implementation) determines the upper limit of pixel values. Pixel values higher than the threshold indicate common areas of all vessels in a particular category.

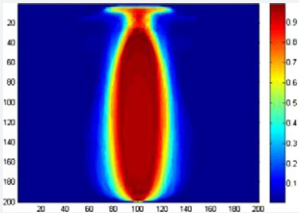
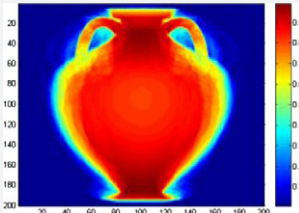
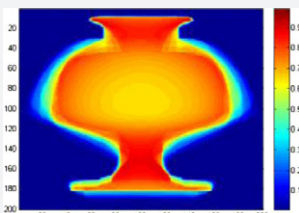
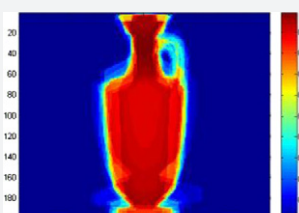
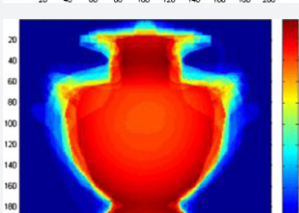
5. Description of the prototype system

We developed a prototype system composed by a web-based search engine, a 3D vessel native-XML database and a stand alone software application. More specifically, the search engine can be found at <http://www.ipet.gr/3DPSE> and allows the user to perform

queries-by-example, to browse by shape category or to perform queries based on specific metadata. The implementation is based on technologies such as PHP [31] and EXIST [32]. On the other hand, the stand alone application was built in Borland Delphi using the GLScene [33] and the VTK [34] libraries. The application is used to perform the mesh preprocessing and the VMBA descriptor extraction. An implementation of the minimum bounding sphere capable of handling a large number of vertices has been used [35]. On the other hand, an initial mesh decimation step allows the software to handle high density digitised vessels [36]. The application can visualise the preprocessing operation in real time and it can also be used as a stand alone tool for studying pottery as it provides information such as the diameter of the body or the appendages of a vessel at any contouring level. The total number of contouring levels for our experiments was set at 32. This was selected as it results a satisfactory distribution of contouring levels capable of capturing the detailed features of the database models. Additionally, the software automatically encodes the VMBA into XML using a predefined schema. We have used the developed application to extract the VMBA descriptor from all models in our database. The average run-time for the preprocessing (operation visualisation mode enabled) and the descriptor extraction phase was 25.6 seconds on an AMD Athlon at 2.2Ghz running Microsoft Windows XP.

The axis of symmetry was successfully recognised for all 1012 vessels in our repository. The current approach failed in detecting correctly the top of the vessels that were reconstructed using the *shape-from-silhouette* 3D scanning technique. This technique is known to fail in capturing concave and hollow areas which are not visible on the object's silhouette. Thus, the vessel's rim appears to be a flat surface and the algorithm fails to detect its top. However, most of the commercial software based on the *shape-from-silhouette* technique provide tools for manual editing of such surfaces to overcome this problem [37]. Additionally, the current implementation might also fail with objects like shallow plates due to the quantisation of 3D space by the limited number of contouring levels (e.g. the model is positioned between two contouring levels and thus cannot be recognised). The circular regression that is used enhances the robustness of the feature extraction and the descriptor remains unaffected by the digitisation noise (vertex coordinate displacement) and the possible cracks and defects found in real artefacts. Nevertheless, the feature extraction is affected by the missing parts of a vessel which result in a minimum bounding sphere different than the one that would have been calculated if the missing parts were present. In a case where half of the vessel's main body is missing, then the remaining vessel will be scaled to fit within the unit bounding sphere and thus the software will detect greater radius values. Additionally, in cases where the top

Table 1
Intraclass shape variation.

Shape name (objects)	Shape variation	Class visualisation
Alabastron (55)	0.6320	
Amphora (96)	0.5036	
Psykter (48)	0.8694	
Lekythos (16)	0.4471	
Hydria (26)	0.4579	

and bottom parts of a vessel are missing the algorithm will also fail to detect the top of the vessel and then depending on the initial orientation of the vessel, the profile radius values might appear vertically flipped. If such cases, we accept the top-bottom orientation given initially by the principal component analysis. Additionally, in our implementation, a vertically flipped version of the descriptor is generated in real time during the shape matching procedure.

6. Content-based retrieval system performance evaluation

The performance evaluation approach that has been followed was based on performing multiple queries using vessels from the five selected categories of Ancient Greek shapes. All selected vessels were used as a query object. In this section, we discuss the similarity metric that has been used and we present the system's performance by using the VMBA and the MPEG-7 SSD [38] descriptors. The performance evaluation metrics that have been used, are those proposed by the annual *3D Shape retrieval contest* [39].

6.1. Similarity metric selection

In order to identify the most appropriate similarity metric for shape matching, we carried out multiple query tests on the database. For the PFV vector matching, we evaluated the performance of six similarity metrics. We questioned the performance of the Euclidean distance, the normalised and weighted Euclidean distance, the Manhattan distance, the earth mover's distance and the Minkowsky distance. The Euclidean distance has been selected due its superior performance and its low computational cost. The next best was the normalised Euclidean. The performance of the earth mover's distance proved to be inappropriate for the PFV vector. On the other hand, we used the Hamming distance as the similarity metric of the AFV. In order to overcome the inability of the metric to handle bit shifting between neighbouring sectors, we introduced the idea of fuzziness in the shape matching procedure. Thus, before applying the metric, we performed a binary dilation on the AFV using a 3×3 square structuring element. The small dimensions of the AFV array used in our experiments allowed




Fig. 11. Average precision – recall graph of VMBA and MPEG-7 3D shape spectrum over the selected vessel shape categories.




Fig. 12. Average F-Score performance of vessel main body and appendages and MPEG-7 3D shape spectrum over the selected vessel shape categories.

the performance of dilation in real time during the shape matching procedure.

The similarity errors (*Euclidean distance error* and *Hamming distance error*) between two vessels are normalised to unit using the maximum possible masses (P_{max} and AF_{max}) of each case. In the case of PFV, P_{max} is equal to the number of contouring levels multiplied by the unit sphere radius while in the case of AFV, AF_{max} is equal

to the number of elements of the 2D array. Additionally, different weights (P_{weight} and AF_{weight}) can be used for varying the importance of the vessel's main body or appendage parts in the shape matching procedure. The total dissimilarity error between two vessels is given by $d(a, b) = \left(P_{weight} \frac{Euclidean_{a,b}}{P_{max}} + AF_{weight} \frac{Hamming_{a,b}}{AF_{max}} \right)$. In our performance evaluation experiments, P_{weight} and AF_{weight} were set to one.

Table 2
Average values of performance scalars.

Shape	1 st Tier		2 nd Tier		NN		AP		E-measure (64 nd)		CG (64 nd)		DCG (64 nd)	
	VMBA	SSD	VMBA	SSD	VMBA	SSD	VMBA	SSD	VMBA	SSD	VMBA	SSD	VMBA	SSD
Alabastron	47.80	47.40	36.06	28.16	0.87	0.83	0.46	0.43	0.51	0.54	28.56	27.32	7.49	7.33
Amphora	78.02	26.76	45.54	19.04	0.84	0.78	0.75	0.25	0.33	0.73	53.10	21.30	12.45	5.46
Hydria	46.15	17.75	24.40	14.86	0.88	0.53	0.33	0.13	0.71	0.80	12.96	8.69	4.32	2.07
Lekythos	55.07	21.48	32.22	14.84	0.93	0.43	0.43	0.11	0.70	0.84	11.93	6.18	3.83	1.61
Psykter	88.97	49.73	49.76	36.11	1.00	0.75	0.86	0.46	0.17	0.49	46.37	28.33	11.80	7.26
Averages	63.20	32.62	37.60	22.60	0.90	0.66	0.57	0.27	0.48	0.68	30.58	18.36	7.98	4.74

VMBA: vessel main body and appendages; SSD: 3D shape spectrum.



Fig. 13. Queries results using the vessel main body and appendages descriptor.

Furthermore, we used the Euclidean distance for comparing the MPEG-7 SSD histograms since it performed equally well. We have also evaluated the performance of MPEG-7 SSD using 100 bins and 64 bins histograms. The 64 bins histograms proved to have superior performance [38]. Thus, all figures corresponding to evaluation results depict the performance of 64 bins histograms. The MPEG-7 Experimentation Model (*XM version 6.1*) was used for the descriptor extraction. It should be mentioned that additional mesh processing was required for several vessels as they were not manifold and the MPEG-7 SSD descriptor extraction process was failing.

6.2. Performance evaluation results

The following graphs and tables depict the overall performance of the VMBA and the MPEG-7 SSD descriptors for the selected vessel categories. The *average precision–recall* performance (Fig. 11) and *average f-score* (Fig. 12) of VMBA always remain higher. This indicates that when the system is using the VMBA descriptor, similar vessels will be ranked higher and thus appear earlier to the user. According to Fig. 12, VMBA achieved a 20% gain within the first hundred retrieved vessels. This shows an increase in the relevant objects that appear within the first hundred ranking positions. For all performance scalars [39] which appear in Table 2, a higher value indicates a better performance. The *E-measure* is an exception to this rule as lower values indicate better behaviour of the system. The nearest neighbour (NN) scalar is in all cases higher than 80% indicating a high probability of retrieving a similar object at the first ranking position when VMBA is used. The results indicate that the VMBA significantly outperforms the MPEG-7 SSD descriptor in all our tests. Fig. 13 illustrates the results of four queries using the VMBA descriptor. The first vessel in each row represents the query object.

7. Conclusions

Digital images, 3D models and multimedia databases of archaeological artefacts play an important role in modern archaeological research projects. In this work, we presented a novel pose normalisation recursive algorithm applicable to complete or nearly complete 3D vessel replicas. This is done by exploiting the most important attribute of vessels, their axially symmetric shape. Based on the features extracted from the pose normalised vessels, we derived a shape descriptor. The extracted features can also be used for the segmentation of a vessel's main body from its appendages or for the generation of other pottery specific descriptors. To evaluate the performance of the descriptor, we constructed a calibrated 3D vessel database that contains digitised, manually modelled and computer generated vessels. The proposed descriptor was compared to the general purpose MPEG-7 SSD descriptor and as was expected, it performed better in all cases. A web-based prototype search engine (<http://www.ipet.gr/3DPSE>) where the user can perform queries-by-example was also developed.

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 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, Dr. Horn-yeu Shiaw, Prof. Robert Jacob and Prof. Gregory R. Crane from Tufts University, the Centre of Advanced Spatial

Technologies of the Hampson Museum, Carlos Hernández Esteban and Francis Schmitt from Télécom Paris for providing us with copies of their digitised 3D vessel collections.

References

- [1] B. A. Sparkes, Greek pottery: an introduction, Manchester University Press, ISBN-10:0719022363, 1991.
- [2] B. A. Sparkes, The Red and the Black: studies in Greek pottery, ISBN 10:0415126614, Routledge, 1996.
- [3] R. M. Cook, Greek painted pottery, 3rd ed., ISBN-10: 0415138604, Routledge, 1997.
- [4] A. Gilboa, A. Karasik, I. Sharon, U. Smilansky, Towards computerized typology and classification of ceramics, *Journal Archaeological Science* 31 (2004) 681–694.
- [5] H.Y. Shiaw, R.J.K. Jacob, G. R. Crane, The 3D Vase Museum: A new approach to context in a digital library, *Proceedings of the ACM/IEEE Conference on Digital Libraries, USA, June 7–11, 2004*, p. 125–134.
- [6] A. Payne, K. Cole, R. Mainfort, C. Goodmaster, K. Simon, S. Smallwood, Fred Limp, The Hampson virtual museum, Center for Advanced Spatial Technologies, Fayetteville, Arkansas, <http://hampsonmuseum.cast.ark.edu>, last accessed at 14/06/2009.
- [7] D. Gorisse, M. Cord, M. Jordan, S. Philipp-Foliguet, F. Precioso, 3D content-based retrieval in artwork databases, in: *Proceedings of the 3DTV-Conference, Kos Island, Greece, May 7–9, 2007*, pp. 1–4.
- [8] C. Hörr, G. Brunnett, Similarity estimation on ancient vessels, *Proceedings of the GraphiCon Conference, Moscow State University, Russia, June 23–27, 2008*, p. 94–100.
- [9] E. Paquet, C. Lahanier, D. Pitzalis, G. Aitken, S. Peters, H.L. Viktor, Content-based indexing and retrieval of cultural heritage data: An integrated approach to documentation with application to the EROS Database, *Proceedings of the 7th International Symposium on virtual reality, archaeology and cultural heritage, Nicosia, Cyprus, October 30–November 4, 2006*, p. 270–275.
- [10] A. O. Shepard, *Ceramics for the archaeologist*, Carnegie Institution of Washington, 1980, ISBN-10:0872796205, p. 225–251.
- [11] The Beazley Archive, <http://www.beazley.ox.ac.uk/index.htm>, last accessed at 6/08/2009.
- [12] C. Maiza, V. Gaildrat, *SemanticArcheo: A symbolic approach of pottery classification*, in: *Proceedings of the 7th International Symposium on virtual reality, Nicosia, Cyprus, October 30–November 4, 2006*, pp. 227–233.
- [13] C. Maiza, V. Gaildrat, Automatic classification of archaeological potsherds, in: *Proceedings of the 8th International Conference on computer graphics and artificial intelligence, Limoges, France, May 11–12, 2005*, pp. 135–147.
- [14] L. De-zhi, R. Anshuman, S. Arleyn, B. Myungsoo, An XML-based information model for archaeological pottery, *Journal Zhejiang University Science* 6A (5) (2005) 447–453.
- [15] L. Dezhi, A. Razdan, Knowledge-based search engine for specific 3D models, in: *Proceedings of the International Symposium on computational and information science, Shanghai, China, December 16–18, 2006*, pp. 530–537.
- [16] Q. Huang, S. Flory, N. Gelfand, M. Hofer, H. Pottman, Reassembling fractured objects by geometric matching, *ACM Transactions Graphics* 25 (Issue 3) (2006) 569–578.
- [17] R. Sablatnig, M. Kampel, Model-based registration of front- and backviews, *Computer Vision Image Understanding* 87 (2002) 90–103.
- [18] M. Kampel, H. Mara, R. Sablatnig, Robust 3D Reconstruction of Archaeological Pottery based on Concentric Circular Rills, *Proceedings of the 6th International workshop on image analysis for multimedia interactive services, Montreux, Switzerland, April 13–15, 2005*, p. 14–20.
- [19] F.J. Melero, A. Leon, F. Contreras, J.C. Torres, A new system for interactive vessel reconstruction and drawing, in: *Proceedings of Computer Applications in Archaeology, Vienna, Austria, April 8–12, 2003*, pp. 8–12.
- [20] Y. Cao, D. Mumford, Geometric structure estimation of axially symmetric pots from small fragments, in: *Proceedings of the signal processing, pattern recognition and applications, IASTED, Crete, Greece, June 25–28, 2002*, pp. 92–97.
- [21] C. Orton, P. Tyers, A. Vince, *Pottery in archaeology*, ISBN-10: 0521445973, Cambridge University Press, 1993, p. 152–163.
- [22] E. Welzl, Smallest enclosing disks, *Lecture Notes Computer Science* 555 (1991) 359–370.
- [23] W.J. Schroeder, J.A. Zarge, W.E. Lorensen, Decimation of triangle meshes, in: *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques, Chicago, USA, July 27–31, 1992*, pp. 65–70.
- [24] R. Israel, Best circle fitting, <http://hdebruin.soo.dto.tudelft.nl/jaar2006/kromming.pdf>, last accessed at 7/08/2009.
- [25] A. Koutsoudis, G. Pavlidis, F. Arnaoutoglou, D. Tsiafakis, C. Chamzas, qp: A tool for generating 3D models of Ancient Greek pottery, *Journal Cultural Heritage* 10 (2009) 281–295.
- [26] PHP Server-side HTML embedded scripting language, <http://www.php.net>, last accessed at 7/09/2009.
- [27] eXist is an open source repository and retrieval engine, <http://exist.sourceforge.net>, last accessed at 7/08/2009.
- [28] GLScene, an OpenGL based 3D library for Delphi, <http://glscene.sourceforge.net>, last accessed at 7/09/2009.
- [29] The Visualization Toolkit (VTK), <http://www.vtk.org/>, last accessed at 7/09/2009.

- [35] B. Gartner, Fast and robust enclosing balls, *Lecture Notes Computer Science* 1643 (1999) 325–338.
- [36] M. Knapp, Mesh Decimation using VTK, http://www.cg.tuwien.ac.at/courses/Seminar/SS2002/Knapp_paper.pdf.
- [37] [3D] Software object modeller Pro, <http://www.3dsom.com>, last accessed at 6/08/2009.
- [38] C. Grana, R. Cucchiara, Performance of the MPEG-7 shape spectrum descriptor for 3D objects retrieval, in: *Proceedings of the Italian Research Conference on Digital Library Management Systems*, Padova, Italy, January 27, 2006, pp. 11–14.
- [39] [3D] Shape retrieval contest, <http://www.aimatshape.net/event/SHREC>, last accessed at 6/80/2009.