

# Developing a Recognition System for the Retrieval of Archaeological 3D Models

Jiménez-Badillo, D.<sup>1</sup>, Ruíz-Correa, S.<sup>2</sup>, García-Alfaro, W.<sup>2</sup>

<sup>1</sup>Instituto Nacional de Antropología e Historia, México

<sup>2</sup>Centro de Investigación en Matemáticas, Guanajuato, México  
*diego.jimenez61@gmail.com, src@cimat.mx, wgarcia@cimat.mx*

---

*This paper describes an ongoing research focused on developing an automatic recognition system for the retrieval of archaeological 3D models. The aim of the project is to implement computer tools capable of exploring a database and retrieving those objects that are similar in shape to a given 3D 'query model'. Such tools would be extremely useful to archaeologists who need to compare shape features between artefacts as part of their quantitative analysis. One interesting case is the comparison of stylistic similarities in anthropomorphic objects recovered from the remains of the ancient city of Tenochtitlan (today's Mexico City).*

*Keywords:* 3D shape retrieval, shape matching, shape recognition.

---

## 1. Introduction

Thanks to recent advances in scanning technologies there has been an increase in the number of methods developed for digitizing three dimensional shapes. Some of these are routinely applied to archaeological artefacts producing vast amounts of valuable information (LEVOY *et al.*, 2000).

The most common products of such efforts are 3D models structured as surface meshes (Figure 1). A surface mesh consists of a collection of vertices, edges and triangular faces that approximate the shape of the real object.

Unfortunately, once a collection of 3D models is stored in a computer repository, it is difficult to gain access to such data due to the lack of appropriate retrieval mechanisms. The conventional way to search for 3D models consists in composing a query based on text descriptions. However, textual annotations are necessarily constrained by the database application domain, as well as by language and other factors. Consequently they are inadequate for general purpose searches.

In order to overcome such limitations and unleash the potential of 3D data in archaeological research, it is worthwhile to explore ways of building an automatic recognition system that searches, retrieves and classifies 3D models by comparing geometric features instead of text.

This paper describes an ongoing research focused on that goal with a special attention to the specific requirements of archaeological applications.

## 2. Main goal

The main goal is to produce a system that takes a 3D surface mesh (i.e. the *query model*) as input, and then compares it to hundreds, possibly thousands of 3D scanned objects in order to identify those that approximate the shape of the query model. The construction of this type of system has been the subject of intense research since the 1980's, especially in fields such as computer vision and geometric modelling (BESL and JAIN, 1985; LONCARIC, 1998; CAMPBELL and FLYNN, 2001).

A paper surveying recent proposals identifies the most common operations involved in the process of matching and retrieving 3D shapes (TANGELDER and VELTKAMP, 2004):

Firstly, each 3D model must be identified by a *shape descriptor*. This is a measure of some geometric attribute -such as curvature, that characterizes the particular form of each object. By analyzing patterns resulting from the aggregation and contrast of shape descriptor values, we are able to obtain the so-called *shape features*, which encode local topological and geometric properties of the model. In turn, the analysis of shape features allows us to discriminate differences between 3D models.



**Figure 1:** A sample of 3D models obtained from 'Mezcala masks' found in the Sacred Precinct of Tenochtitlan. The surface meshes are rendered as solids in false colour.

Secondly, shape features must be stored within an index structure to facilitate search and matching operations.

Thirdly, a special algorithm must be put in place for calculating a dissimilarity distance between a particular query model and the rest of the objects in the database. Matching similar objects becomes a task of comparing the shape descriptors stored in the index structure.

Fourthly, a fetching algorithm provides functionality such as the ranking of objects according to their level of similarity.

Finally, the fifth operation allows for efficient visualization and/or interaction with the user.

In the following pages we describe our efforts in implementing some of these operations, while advancing some ideas on how to apply this to an archaeological collection.

### 3. Requirements

Any recognition system must comply with requirements such as efficiency and robustness, but in the case of archaeological applications it is necessary to satisfy extra demands. The first one is to be able to discover shape similarities between objects that are partially isometric, that is, objects that are not identical but share at least some common features. Partial matching is the task of finding similar sub-parts among objects, and it is a harder problem to solve than measuring similarity between entire models (i.e. global matching) (GAL and COHEN-OR, 2006).

Partial matching is particularly important in archaeological situations where not all artefacts are complete. One may need to retrieve, for example, all instances of anthropomorphic figurines, without being affected by the fact that one of them lacks the head, while another is missing one leg. Partial matching is also useful to retrieve complete models of the same class, regardless if they differ in some details (Figure 2).

A second requirement is for the system to be able to detect similarities without being affected by model variations in rotation, translation, reflection and scale. Ideally, an arbitrary combination of translation, rotation and scale applied to one object should not affect its similarity measure with respect to another object.

In summary, a specialized set of shape analysis methods is required. The challenge is to apply algorithms which identify and rank shape similarities efficiently, within the limits of current computer power and discriminative enough to perform partial matching.



**Figure 2:** Sculptures belonging to the same class of seated Aztec deities. Both should be retrieved during a general-purpose query despite partial differences in head and body features.

### 4. Approach

There are two classes of shape descriptors: global and local. As the term implies, global descriptors focus on measuring properties belonging to the entire 3D model, such as the statistical moments of the boundary of the model, volume-to-surface ratio, and the Fourier transform of the boundary of the model (check more examples in TANGELDER and VELTKAMP, 2004).

Alternately, the computation of local descriptors relies on selecting a certain number of points on the surface of the model in order to measure curvature or other properties within the neighbourhood of each selected point. The limits of such neighbourhood are commonly defined by a *spherical kernel*.

There are many approaches to obtain local shape descriptors. One is to compute shape index values over an entire mesh to produce a histogram that describes the local features of the 3D model (ZAHARIA and

PRETEUX, 2001). Another method is to compute *point signatures* that reveal surface information along a three-dimensional curve in the neighbourhood of points (CHUA and JARVIS, 1997). A third approach is to define patches around each point and then estimate curvature by fitting a quadric patch that approximates the surface as best as possible (GAL and COHEN-OR, 2006).

The approach followed in this project is based on extracting local shape descriptors by measuring principal curvatures in local neighbourhoods defined at different scales. Analysing the aggregation and contrast of shape descriptors enable the extraction of shape features, which encode the topology and local geometry of the 3D model. In summary, our algorithm performs the following steps:


For each triangular surface mesh:

1. Extract local shape descriptors by estimating the so-called *principal curvature* values. This task is divided into two steps. The first one consists in defining a region around each point using a spherical kernel. Such kernel region constitutes a local neighbourhood of a particular point. By using kernels of different sizes it is possible to define local neighbourhoods at different scales, making the algorithm multi-scale, multi-resolution. Afterwards, during the second step, we quantify how much the surface bends in different directions from the centre point of each local neighbourhood. This is done by computing two principal curvature measures (i.e.  $k_1$  and  $k_2$ ), as well as two derived measures corresponding to mean curvature =  $(k_1+k_2)/2$ , and Gaussian curvature =  $(k_1.k_2)$  (DO CARMO 1976).
2. Project the 3D mesh of each object into a 2D embedding space using a geometric transformation called locally linear embedding (LLE), proposed by ROWEIS and SAUL (2000). This step is extremely important because enable us to compare local geometric and topological properties between 3D models.
3. Calculate how much *bending energy* is necessary to transform the 2D LLE embedding of the query model into the 2D LLE embedding of the rest of the models.
4. Create an index structure with the values extracted in the previous step.
5. The fetching and visualization tools are still under construction and we will report these in a future paper.

#### 4.1. Estimate curvature to extract shape descriptors

As mentioned earlier, the first step of this project is to extract local shape descriptors by estimating four curvature values. To understand the process, consider


the simple situation represented in Figure 3, in which the level of bending of a 1D surface (i.e. the curved line) can be estimated by means of a straightforward procedure. We would simply have to fit a ball to the curve and then calculate the inverse function of its radius with the formula  $c=1/r$ ; where  $r$  is the radius.



**Figure 3:** Fitting a ball to measure curvature in an ideal situation:  $c = 1/r$ ; where  $r$  is the radius.

Unfortunately, estimating curvature in triangular surface meshes is a more difficult problem to solve, mainly because curvature shows distinct profiles along different directions.


In addition, triangular surface meshes representing real objects contain unpredictable levels of noise in the form of outlier points. Figure 4, for example, shows the difference between an ideal, expected curve (dotted line) and the real, imperfect profile produced by the scanning process (solid line). It is clear that many points of the real mesh lie outside the ideal expected curve, producing a noisy zigzag profile. Fitting a ball to calculate the curvature around point  $p'$  –which has an outlying neighbour–, would lead to inaccurate measurements.



**Figure 4:** Real profile on the surface mesh of a 3D model. The dotted line represents the ideal shape, while the solid line corresponds to the real profile captured by a scanner.

Consequently, such non-robust procedure of curvature estimation is too sensitive to the quality of the mesh, to its resolution, and to the smoothness/roughness of the original surface. When applied to a real triangular sur-

face mesh, like the one shown in Figure 5, the method fails to reveal any significant shape feature that can be used later to compare this mask with other 3D models.



**Figure 5:** Application of non-robust curvature measurement. Shades of colour correspond to curvature values for each local neighbourhood on the surface of this mask. Notice that extreme values (blue and red) are uniformly distributed and therefore no identifiable features could be found during a matching process.

To avoid that problem, we apply an alternative, more robust method. This consists in estimating the principal curvatures -along principal directions- of several local neighbourhoods around each point of the mesh. Such local neighbourhoods are defined by kernels of different sizes. In this way, the estimation of curvature becomes multi-scale, multi-resolution.

As explained by one of the authors of this approach, the method works well for polyhedral approximations with a large number of small faces, such as a typical triangular surface mesh (TAUBIN, 1995). More importantly, the process allows for the identification of shape descriptors at different scales and with different levels of detail.

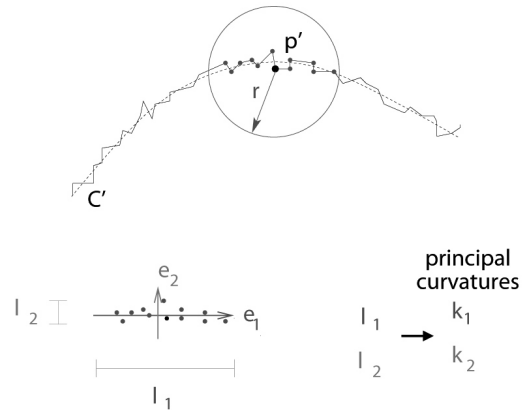
The two principal curvatures at a given point of a surface mesh are the maximum and minimum values of the curvature. They measure how the surface bends by different amounts in different directions at that point. They are denoted  $k_1$  (i.e. maximum curvature) and  $k_2$  (i.e. minimum curvature) and correspond to the Eigenvalues of the extrinsic curve at that point (Figure 6). Eigenvalues are calculated through a standard application of principal components analysis (PCA) using an algorithm developed by YANG *et al.* (2006).

Additionally, we derive the median and Gaussian curvatures from the previously computed principal curvatures using the formulae:

$$\text{Median curvature} = (k_1+k_2)/2$$

$$\text{Gaussian curvature} = k_1.k_2$$

Once the four curvature measures are estimated, their distribution is statistically analysed in order to discover patterns that correspond to perceptually meaningful *local shape features*. Local shape features are, therefore, aggregations of shape descriptor values.




**Figure 6:** Principal curvature measurement. Notice the kernel delimiting a local neighbourhood and the extraction of Eigenvalues to measure how much the surface bends in the horizontal and vertical directions.

A local shape feature is considered ‘salient’ (and potentially significant) if the curvatures of the local neighbourhoods result highly unpredictable with respect to adjacent neighbourhoods after testing it at different scales. Unpredictability is determined as a function of the local probability density function computed on every local neighbourhood of the point. Such definition of saliency is inspired by the work of KADIR and BRADY (2002).

Such features are frequently, though not exclusively, characterized by contrasts between curvature values, in other words, sharp differences between convex and concave ‘creases’ or ridges and ravines on the surface of the model.

To illustrate the results from this task, we represent curvature values as different shades of colour (Figure 7). Notice how the image on the top left shows no significant features. This is not surprising because it corresponds to the discarded non-robust method of curvature estimation. In contrast, the other three images illustrate results from the more robust method that we propose. They correspond to  $k_1$ ,  $k_3$  and  $k_4$  ( $k_2$  is not included due to lack of space). It is clear that these measures reveal meaningful local shape features, especially in those areas corresponding to the eyebrows, eyes, mouth and nose. The planarity in the remaining parts of the mask is also well extracted by the algorithm as a relatively uniform shade.

To demonstrate the multi-resolution capacity of the method, Figure 8 shows principal curvature estimation for a figurine using kernels of different sizes. Local shape features detected at different scales and resolutions are used later –during the matching phase- to compare a particular surface mesh with the rest of the 3D models.




**Figure 7:** Comparison of results from the non-robust method (top-left) against three principal curvature measures. The results are ordered as follows: top-right maximum curvature ( $k_1$ ); down-left mean-curvature ( $k_3$ ); and down-right Gaussian curvature ( $k_4$ ). The last three are more effective to identify significant shape features, such as the eyebrows, eyes, mouth and nose, as well as the planar surface of the remaining parts of this mask. In contrast, the first image shows too much noise.

Finally, it is worth mentioning that the approach used for computing curvatures performs well regardless the density of the mesh. This is very convenient when one is trying to compare a coarse surface mesh against another mesh possessing higher resolution.


#### 4.2. Mapping surface meshes into 2D embedding space

The next step is to project the surface mesh of each object into a 2D embedding space. The strategy resembles the projection of geographic spherical coordinates of the Earth into a plane map. Both in the case of cartographic mapping and 3D models it is necessary to transform object descriptions from 3D to 2D space but preserving important geometric properties, such as shape.

For the purposes of this project, the so-called *locally linear embedding* (LLE) is considered an appropriate methodology for performing the mapping (ROWEIS and SAUL, 2000; SAUL and ROWEIS, 2003; RIDDER and



**Figure 8:** The effect of applying kernels of different sizes in order to reveal features at different scales. The kernel radius is determined as a percentage of the longest distance from the mass center of the object to its surface. In this example, the kernels are 10%, 15% and 25%.



**Figure 9:** A 3D model (right) embedded into 2D space (left) by means of the locally linear transformation. The local shape features delimited by shades of colour (i.e. curvature measures) are preserved after the transformation.

(DUIN, 2002). This approach is based on the geometric intuition that each surface mesh point and its neighbours are expected to lie on or close to a locally linear patch. This observation holds both on the surface of the original 3D space and on the surface of the embedding 2D space. This means that both 3D geometric information as well as local topological properties of the original surface mesh are approximately preserved in the embedding 2D space. Therefore, it is assumed that the 2D mapping provides enough information for the purpose of comparing the objects in search for similarity. Equally important is the fact that LLE allows reconstructing the original 3D model from its 2D embedding.

Figure 9, for example, shows the appearance of a mesh before and after been subjected to a locally linear embedding. The shades of colour correspond to curvature values and are included in the image to show how the local shape features existent in 3D are preserved in the 2D embedding.

The advantage of this strategy is that assessing similarities between shapes becomes much easier when the surface meshes are projected into 2D, mainly because the reduction in dimensionality facilitates the comparison of geometric and local topological properties, as explained in section 4.3.

#### 4.3. Comparing 2D surfaces for matching

The third step seeks to find a transformation that deforms the 2D LLE projection of each surface mesh store in the database into the 2D LLE projection of the query model. Such transformation is called *global deformation mapping*.

The process of finding the global deformation mapping is aided by taking into consideration the previously computed values of curvature, whose statistical variation patterns throw light on the most significant parts of the model. These are used as landmarks with which comparisons between models can be done.

The transformation must be such that landmark features from one surface mesh should approximately match similar landmark features in the query surface mesh. In order to assess similarity between the landmark features of the two meshes, an analysis of local neighbourhoods is performed using kernels at different scales as proposed by (SCHOLKOPF and SMOLA, 2002).


The global deformation mapping is computed applying numerical analysis (NOCEDAL and WRIGHT, 2006). The process finds the mapping that minimizes the transformation mean square error between surface meshes. The output is an error term that quantifies how much “energy” is required to deform one surface to fit it into another. Similar shapes lead to smaller energy deformation errors than shapes that are distinctively different. It is precisely this error term which is stored in an index structure to enable the ranking of objects by degree of similarity during the fetching stage. As we mentioned above, the fetching and visualization algorithms are still work in progress.

## 5. Study case

The idea for this project came from the need to rank similarities in hundreds of anthropomorphic stone figurines and masks found in the remains of the Sacred Precinct of Tenochtitlan, the main ceremonial Aztec complex, located in Mexico City (Figure 1).

The schematic features of these objects set them apart from other artefacts with more naturalistic style. This has attracted the attention of many specialists and during the last three decades these items have been the subject of intense debate for two main reasons:

First, the 220 figurines and 162 masks were located in 14 Aztec offerings dating from 1390 to 1469 A.C., yet they do not show typical “Aztec features”. Indeed, their appearance resembles artefacts from the southern State of Guerrero, particularly from the Mezcala region, which is hundreds of kilometres away from the ancient Tenochtitlan. Such origin would not be rare, as it was common for the Aztecs to import goods from other regions either by trade or by extracting tribute from conquered towns. The style of the masks and figures, however, is more difficult to explain. It is similar, if not identical, to the style of objects manufactured in Mezcala and other places of Guerrero during much earlier times, probably during Classic (200-1000 A.C.) or even Preclassic times (2000 B.C. to 200 A.C.), while the offerings correspond to a Late Postclassic contexts (1340-1521 A.C.). This leads to the question: Did the Aztecs collect ‘antique’ objects to re-use them in their own offerings? Or did the Guerrero/Mezcala styles survive till the late Postclassic period and therefore the offerings objects were produced during Aztec times? It is worth noticing that before the finding of these Aztec offerings very few Mezcala style artefacts were found in Postclassic contexts. Unfortunately, not enough stratigraphic information is available for the collections from



**Figure 10:** *Typology of nose-shapes taken from the Mezcala collection studied by OLMEDO and GONZALEZ (1986). Recognizing such features would be easier by using an automatic 3D shape recognition system like the one we propose here.*

Guerrero, so specialists rely purely on stylistic considerations to explain the affiliation and chronology of these artefacts.

Second, it is not clear how many Guerrero/Mezcala styles exist. Some specialists believe there are at least five different traditions (COVARRUBIAS, 1948, 1961; OLMEDO and GONZALEZ, 1986; GONZALEZ and OLMEDO, 1990), while others recognize only four (GAY, 1967) or two (SERRA PUCHE, 1975). The diversity of views is due, in part, to a lack of contextual information available for the majority of artefacts found in Guerrero, but it also reflects the subjective criteria used to classify such artefacts.

Clearly, more objective methods are needed to answer questions such as: how many styles were developed in the Guerrero/Mezcala regions? How many of these styles coexisted? Were some styles contemporary with the Aztecs? Which specific styles are represented among the 382 offering objects found in the Sacred Precinct of Tenochtitlan?

Previous studies have tried to solve some of these questions by classifying object shapes with heuristic methods. In one of these studies, due to OLMEDO and GONZALEZ (1986), the forms of faces, eyes, noses, etc. are used to rank object similarities in order to define several object types (Figure 10). We believe that a more objective classification could be achieved by further expanding the functionality of our matching tools in order to compare salient features discovered automatically by our recognition system. As illustrated above (Figure 7) our method is able to distinguish mask features (eyebrows, nose, eyes, mouth, etc.) relevant for this kind of application.

## Conclusion

Computer tools for shape matching and retrieval that are designed specifically for archaeological research could provide new capabilities for classifying correspondences among artefact collections and may help to resolve some stylistic questions. We are currently testing our algorithms with the Mezcala collection and the results are encouraging. As the tools developed in this project are generic, we expect that they would prove useful in the analysis of other archaeological collections.

Especially important is the capacity of our system to perform partial matching through a) a robust method of curvature estimation and b) the comparison of local topological and geometrical features using the so-called local linear embedding. Once the remaining modules of our system are completed, we expect to perform a benchmark analysis, whose results would be published.

## References

- BESL P. J., JAIN R. C. 1985. Three-dimensional object recognition. *Computing Surveys* 17 (1), pp. 75-145.
- CAMPBELL R. J., FLYNN P. J. 2001. A survey of free-form object representation and recognition techniques. *Computer Vision and Image Understanding* 81 (2), pp. 166-210.
- CHUA S. J., JARVIS R. 1997. Point signatures: a new representation for 3D object recognition. *International Journal of Computer Vision* 25 (1), pp. 63-65.
- COVARRUBIAS M. 1948. Tipología de la Industria de Piedra Tallada y Pulida de la Cuenca del Río Mezcala. In *El Occidente de México*. Sociedad Mexicana de Antropología, Mexico, pp. 86-90.
- COVARRUBIAS M. 1961. *Arte Indígena de México y Centroamérica*. Universidad Nacional Autónoma de México, Mexico.
- DO CARMO M. P. 1976. *Differential Geometry of Curves and Surfaces*, Prentice-Hall.
- GAL R., COHEN-OR D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Transactions on Graphics*.
- GAY C. T. 1967. *Mezcala Stone Sculpture: The Human Figure*. The Museum of Primitive Art: Studies Number Five. New York.
- GONZALEZ C., OLMEDO B. 1990. *Esculturas Mezcala en el Templo Mayor*. Instituto Nacional de Antropología e Historia, Mexico.
- KADIR T., BRADY M., 2001. Saliency, Scale and Image Description. *International Journal of Computer Vision* 45(2), pp: 83-105.
- LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D. 2000. The digital Michelangelo project: 3D scanning of large statues. In *Proc. SIGGRAPH 2000*, pp. 131-144.
- LIPMAN Y., FUNKHOUSER T., 2009. Möbius Voting for Surface Correspondence. In *ACM Transactions on Graphics: Proc. SIGGRAPH '09* (Aug. 2009), vol. 28 (3).
- LONCARIC S. 1998. A survey of shape analysis techniques. *Pattern Recognition* 31 (8), pp. 983-1001.
- NOCEDAL J., WRIGHT S. J., 2006. *Numerical Optimization*. Springer-Verlag, New York.
- OLMEDO B., GONZALEZ C. 1986. *Presencia del Estilo Mezcala en el Templo Mayor: Una Clasificación de Piezas Antropomorfas*. Thesis presented as requirement for the degree of Bachelor in Archaeology. Escuela Nacional de Antropología e Historia, Mexico.
- POTMAN H., WALLNER J., HUANG Q., YANG Y. L. 2009. Integral Invariants for Robust Geometry Processing. *Computer Aided Geometric Design* 26 (1), pp: 37-60.
- RIDDER D., DUIN R. 2002. Locally linear embedding for classification. *Pattern Recognition Group Technical Report Series*, N. PH-2002-01, 13 pp. (Submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence).
- ROWEIS S., SAUL L., 2000. *Nonlinear Dimensionality Reduction by Locally Linear Embedding*. *Science* 290 (5500), (Dec. 22, 2000), pp. 2323-2326.
- TANGELDER J., VELTKAMP R. C. 2004. A survey of content based 3D shape retrieval methods. In *Proc. International Conference on Shape Modelling and Applications*, Genova, Italy, pp. 145-156.
- TAUBIN G. 1995. Estimating the tensor of curvature of a surface from a polyhedral approximation. *IEEE Proceedings of the Fifth International Conference on Computer Vision (ICCV'95)*, pp. 902-907.
- SAUL L., ROWEIS S. 2003. Think globally, fit locally: Unsupervised learning of low dimensional manifolds. *Journal of Machine Learning Research* 4, pp. 119-155.
- SCHOLKOPF B., SMOLA A., 2002. *Learning with Kernels*. Cambridge University Press, London.
- SERRA PUCHE, M. 1975. Intento de seriación en culturas de Guerrero. Cronología del Estilo Mezcala. In *Proc. XIII Mesa Redonda de la Sociedad Mexicana de Antropología*, Jalapa, Mexico.
- YANG T. L., LAI Y. K., POTMAN H., 2006. Robust principal curvatures on multiple scales. *ACM International Conference Proceeding Series; vol. 256. Proceedings of the fourth Eurographics Symposium on Geometry Processing*, Clagiari, Sardinia, Italy (June 26-28, 2006), pp. 223-226.
- ZAHARIA T., PRETEUX F. 2001. 3D shape-based retrieval within the MPEG-7 framework. *Proceedings of the SPIE Conference 4304*, pp. 133-145.

