Categorization of Aztec Potsherds using 3D Local Descriptors

Edgar Roman-Rangel[1], Diego Jimenez-Badillo[2], Estibaliz Aguayo-Ortiz[2]

[1] CVMLab - University of Geneva, Switzerland
[2] National Anthropology and History Institute of Mexico (INAH), Mexico

Abstract. We introduce the Tepalcatl project, an ongoing bi-disciplinary effort conducted by archaeologists and computer vision researchers, which focuses on developing statistical methods for the automatic categorization of potsherds; more precisely, potsherds from ancient Mexico including the Teotihuacan and Aztec civilizations. We captured 3D models of several potsherds, and annotated them using seven taxonomic criteria appropriate for categorization. Our first task consisted in exploiting the descriptive power of two state-of-the-art 3D descriptors. Then, we evaluated their retrieval and classification performance. Finally, we investigated the effects of dimensionality reduction for categorization of our data. Our results are promising and demonstrate the potential of computer vision techniques for archaeological classification of potsherds.

1 Introduction

The application of computer vision technologies for the preservation, management, and analysis of cultural heritage artifacts has witnessed a rapid growth during the last decade [1]. This is especially true with regard to the creation and use of digital 3D models, which enable capabilities that would not be available using the original artifacts, such as automatic and semi-automatic content analysis [2, 3], virtual reconstructions[4, 5], more efficient archiving [6, 7], sharing documentation online [1, 7], training of novel scholars, etc.

An area of especial interest is the statistical analysis of features observed on 3D models of potsherds with the purpose of categorizing archaeological pottery [3], which is one of the most important jobs in archaeology [2, 4]. In this work, we use the term *categorization* with a slight different connotation than that of *classification*. More precisely, categorization defines a series of decisions that have to be made about a sherd of interest, which is different from the traditional classification scenarios where a single class label is assigned to a query instance. For example, besides of classifying a sherd in function of the type of vessel it comes from, we might also be interested in knowing the type of rim it has. These distinctions are of high relevance, as knowing certain of such characteristics helps archaeologist to infer the final class of the sherd.

This paper introduces the Tepalcatl project, an ongoing effort that brings together archaeologist from the National Anthropology and History Institute of Mexico (INAH) and researchers in computer vision. This project is oriented to developing a system for the automatic categorization of ceramic sherds. Fig. 1, shows examples of Aztec potsherds used as study case in this work.

(a) Cajete. (b) Cajete Tripode. (c) Cup.

Fig. 1: Pictures provided by the project Urban Archaeology of the INAH directed by the archaeologist Raul Barrera. These artifacts are from excavations conducted near the Templo Mayor in Mexico City.

The contributions of this paper are as follows.

1. We design a recognition system for potsherd. First, we acquired a set of 3D surfaces of sherds belonging to specific vessel shapes. These model instances are then used for semi-supervised training of the recognition system.
2. We reformulate the archaeological task of sherd categorization into a computer vision problem, namely 3D image classification and retrieval. For this purpose, we defined seven different taxonomic criteria for which the recognition system is able to perform sherds categorization. For example, it is able to distinguish between several types of rim; or using another criterion, it is able to recognize the type of curvature of the fragment. Note that the training of the system happens only once, but it works well with the seven different criteria nevertheless.
3. We exploit the discriminative power of two state-of-the-art local descriptors, namely the Scale Invariant Feature Transform (SIFT) [8, 9] and the Spin Images descriptor [9, 10]. These descriptors are invariant to rotation and scale transformations, and therefore avoid the need of a pre-processing step to normalize the scale and orientation of the sherds, as required by other methods [3, 11]. Our approach proved to be effective for potsherd categorization.

Our approach allows examining different combinations of local descriptor, dictionary sizes and taxonomic criteria in order to find out which one achieves the best performance. To the best of our knowledge, this approach has not been tried before in archaeological classification.

The rest of this paper is organized as follows. Section 2 describes the details of the Tepalcatl project. Section 3 discusses detailed work in areas of digital preservation of pottery and their automatic analysis. Section 4 summarizes the statistical descriptors that we used. Section 5 explains the data that we used and our experimental protocol. Section 6 discusses our findings. Finally, section 7 presents our conclusions and future directions for our project.

2 Tepalcatl project

The analysis and classification of pottery sherds constitutes one of the most important jobs in archaeology. It provides cultural information of past societies at multiple scales, from the identification of human activity areas to the determination of site chronology and/or the elucidation of regional economic systems.

Unfortunately, a detailed analysis of ceramic fragments is also one of the most cumbersome and time-consuming tasks for archaeologists. This is due not only to the long learning curve involved in mastering a ceramic classification system but also to the vast quantity of sherds normally recovered from the field. A typical excavation in Central Mexico, for example, produces tens of thousands of fragments and a single site may contain ceramics dating from a very long period, encompassing occupations of the Teotihuacan, Toltec and/or Aztec civilizations, that is, sherds dating from approximately 100 B.C.E. to 1521 C.E. Fig. 2 shows the settlement region of the Teotihuacan, Toltec and Aztec civilizations.

(a) Settlement region of Teotihuacan, Toltec, and Aztec cultures.

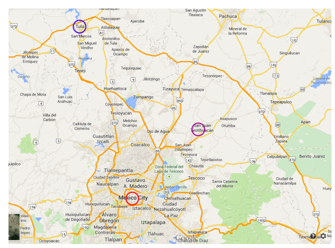

(b) Purple circle: Teotihuacan; blue: Toltec; red: Aztec.

Fig. 2: Settlement region of Teotihuacan, Toltec (city of Tula), and Aztec (city of Tenochtitlan) cultures. (a) Localization within current Mexico. (b) Details of (a) indicating the capital cities.

2.1 Sherd categorization

A very relevant property in ceramic analysis is shape [3]. Archaeologists are expected to reconstruct the profile of a whole vessel by examining the form of a surviving fragment. A traditional method to accomplish this task, has been a visual comparison of a sherd with a manual drawing of the silhouette and diameter of vessels already known. However, this approach is time consuming and requires the knowledge of highly trained people. A common practice to reduce the complexity of the task is selecting only fragments that are considered diagnostic to identify important ceramic types. This includes certain body parts like rims, legs, handles, etc., which are then used in a more detailed typological analysis to deduce chronology, cultural style, etc. Fig. 3 shows examples of potential sherds from the type of ceramic we use.

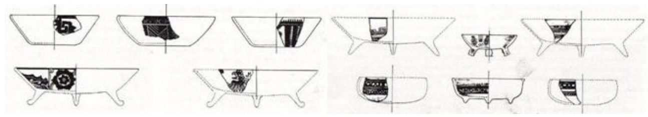

Fig. 3: Potential sherds from the type of pottery pieces in our project.

During the past decade, some professionals have been experimenting with digital technologies oriented to reducing learning curves and improving the quality of the classification process. One of the most interesting proposals has been the acquisition of 3D models of sherds with the purpose of applying mathematical, computer vision and/or machine learning techniques that perform classifications more accurately in an automatic or semi-automatic way [2, 3]. The extraction of 3D digital data has also brought extra benefits, such as the possibility to undertake new types of content analyses, as well as an easier sharing of information among professionals, the design of better ceramic documentation and archiving systems [6, 7], and the performance of virtual reconstruction of vessels [4, 5].

2.2 Goals of the Tepalcatl project

To help in this categorization endeavor, we decided to experiment with new techniques, hoping that a better and more efficient classification method will emerge from successive iterations that benefit from the join efforts of archaeologists and computer vision specialists.

Namely, the goals that our project pursuits are:

1. **Design and compilation of a 3D dataset**. We will digitize a large collection of potsherds from the Aztec and Teotihuacan cultures that ranks in the order of tens of thousands. This task will generate a new digital dataset that poses several challenges in terms of visual description and automatic categorization. This dataset will be rich in terms of labels, as its instances will be annotated using different taxonomic criteria.
2. **Advance the state-of-the-art in 3D classification**. We will develop new methods for better description and representation of 3D models. These new models will overcome limitations potentially found in current approaches. Also, we will explore more efficient methods for classification.
3. **Improve the archaeological process of ceramic classification**. Using different taxonomic criteria (explained in section 5.1), we will conduct several analyses of similarity that could lead us to design efficient methods for accurate categorization of potsherds.
4. **Further assessment**. We will assess the potential that our and other methods that could have to deal with more archaeological tasks, such as retrieval of complete vessels from potsherds, and virtual reconstruction of vessels.

2.3 Work in data collection

Our project contemplates the study of potsherds belonging to the ancient Teotihuacan (100-600 C.E.) and Aztec civilizations (1325-1521 C.E.). Teotihuacan was the first urban center in central Mexico and produced one the most influential cultures in the ancient Americas, whereas the Aztecs were the dominant civilization in the same region by the time of the European arrival. Remains of both cultures, including ceramic sherds, are often found in different strata of the same archaeological sites. Fig. 2 shows the settlement location of these cultures.

Aztec III Black on Orange ceramics. The potsherds used in this work belong to a ceramic style known as "Aztec III Black on Orange", which is one of the most important wares in Mesoamerican archaeology. Along with the so-called Aztec I, II and IV it is part of the sequence that helps understand the cultural evolution in Central Mexico during Postclassic times [13].

The geographic distribution of Aztec III Black on Orange covers the entire Basin of Mexico. Vessels of this type were part of the utilitarian assemblage common in households during the late Aztec period (1350-1520 C.E.). For this reason, it is considered a diagnostic ware to infer the presence of Aztec settlements as opposed to those controlled by previous ethnic groups. Another reason for its importance is that the emergence of this style coincides with the formation and development of the Aztec empire.

Vessels belonging to this ceramic complex show black painted designs drawn on a polished orange surface. The designs are formed by thin brush strokes that follow parallel and curved lines combined with other motifs, such as birds, fishes, plants, geometric figures etc. The designs are found both on the interior and exterior surfaces of the vessels. Fig. 1 shows examples of the Aztec III Black on Orange ceramics. A feature of this ceramic complex, and a major challenge for the classification of sherds, is the abundance of shapes given to the vessels. The most common are cajetes with three legs (Fig. 1b), comales (a flat circular surface to cook corn tortillas), bowls, jars, pots, plates and the so-called apaxtles.

A great quantity of this type of ceramics was collected on the surface during the 1970's by the Valley of Mexico Survey Projects [14–17]. During the 1990's, paste composition and stylistic analyses of those materials was performed, comparing their geographic distribution with historically-documented polities [18, 19]. They found that a number of decorative motifs are exclusive to particular paste groups and managed to identify three main production areas in the Basin of Mexico. This in turn allowed to hypothesize about the exchange systems of the Aztec empire. Thanks to these studies we know in greater detail the economic relations between dependent communities and the Aztec capital.

Most recently, Mexican researchers have uncovered new material inside the boundaries of the Sacred Precinct of Tenochtitlan, the main religious center of the Aztec capital (whose remains lay in close proximity to the main square in Mexico City). The findings have drawn the attention of many specialists, mainly because it proceeds from controlled excavations as opposed to surface surveys.

However, the great quantity of findings has made it difficult to classify and publish the reports.

3 Related work

Computer-aided classification and reconstruction of ceramics has been a subject of research for at least 20 years, with pioneering work by Hall and Lafling [20], who created a software package called NEWTS for drawing, archiving, and editing ceramic profiles using B-spline functions.

Since then, the focus has continued to be the acquisition of 2D profiles, either from sherds or from complete vessels, which can then be classified by archaeologists using their curvature functions [3]. A notable example of these efforts is the creation of the so-called profilograph (`http://www.dolmazon.de/profilograph_e.htm`), a device to draw profiles from 3D objects (i.e. sherds).

Thanks to the growing accessibility of 3D laser scanners, some researchers have shown the advantages of acquiring 3D models of ceramics with the purpose of extracting automatically the 2D profiles [4–6, 21–25]. They based their work on the assumption that the majority of vessels studied by archaeologists are axially symmetric. For perfectly symmetric vessels the profile corresponds to a cross section in the direction of the rotational axis. Thus, by calculating the rotational axis, different methods are able to draw the profile of the vessels with different levels of success. Unfortunately, archaeological vessels are not perfectly symmetric, especially if they were hand-crafted as opposed to thrown-wheel manufactured. Therefore, extracting the rotational axis possesses serious challenges [3]. Perhaps the best method to address this tasks, is the one proposed by Karasik and Smilansky [3] because it takes into account the many possible deviations in the axis of symmetry.

The calculation of the axis of symmetry also requires a pre-processing stage of orientation and alignment of the sherds [26], which is normally a time consuming task [21]. Our analytic approach is invariant to rotation and scale and therefore avoids the need of normalizing orientation and scale.

Another drawback of the current approaches is that they require a whole model of the sherd. This involves taking several scans of the sherd, which are then merged together to complete the model [3]. Our method does not rely on capturing a whole model. Instead it uses only surface information from the back and front views of the sherd, which expedites the process significantly.

As for the comparison of sherds for classification, some methods compare sherds using a point-to-point matching process [27]. Our method compares histograms (i.e. bag-of-visual-words) that represent the statistical frequency of local descriptors (i.e. SIFT and spin images).

Perhaps the most complete and sophisticated approaches under the paradigm of profile extraction that include classification of sherds are due to Gilboa [28] and to Hörr [26]. Hörr [26] proposed a series of mathematical algorithms for profile segmentation, feature extraction and, more importantly, clustering of sherd descriptions using algebraic functions, which in turn allows query and hierarchical classification. The method, however, was applied to complete or semi-complete vessels and results for sherds are not available in the original paper

4 Descriptors

This section describes the two local descriptors that we used in our work. Namely, the Local Invariant Feature Transform (SIFT) [8] and the Spin Images [10], that were originally proposed to deal with 2D images, and whose 3D version consists in scale invariant points of interest in a 3D space [9].

4.1 Spin images

Spin images are well known local descriptors developed to match mesh-surfaces by the individual point-to-point matching of its vertices [10]. However, it might as well be implemented under a bag-of-visual-words approach [9, 12].

In the spin images methodology, each vertex $\boldsymbol{v_i}$ is characterized by

$$\boldsymbol{v_i} = (x_i, y_i, z_i, \theta_i, r_i, h_i)\,, \tag{1}$$

where, (x_i, y_i, z_i) defines the three coordinates position of the vertex; θ_i denotes its local orientation, which corresponds to the orientation of the plane touching the closest points that are connected to $\boldsymbol{v_i}$; and r_i and h_i are respectively the radius and height of a cylindrical support volume centered at (x_i, y_i, z_i).

The base idea of this method consists in *scanning* the 3D surface by *spinning* a sheet around the axis defined by the point orientation θ_i while counting the amount of nearby points [10]. More formally, the cylindrical support volume is divided into R radial segments (*rings*) and L vertical segments (*layers*), thus generating $K = RxL$ spatial bins. Finally, the counting of the number of nearby points $\#\{p_j\}$ falling within each resulting bin $b_k, k = 1, \ldots, K$, is arranged as a K-dimensional feature vector s termed spin image [29],

$$s\,(k) = \#\{p_j : p_j \in b_k\}. \tag{2}$$

The above mentioned methodology can be applied to each vertex in the mesh [10]. However, it might also be applied only to certain vertices selected randomly or by other approaches. One of such approaches is presented in section 4.3.

4.2 SIFT

The Scale-Invariant Feature Transform (SIFT) is a state-of-the-art and one of the most popular image descriptor for gray-scale images. It detects points of interest as points that maximize the difference of Gaussian (DoG) response in a Gaussian scale space, where such a maximization is evaluated both in location and scale spaces. The resulting points are invariant to scale and rotation variations [8].

Recently, the Local-Depth SIFT (LD-SIFT) variant was proposed to deal with 3D images [9], representing the vicinity of an interest point as a depth map. In this approach, the 3D mesh around vertex $\boldsymbol{v_i^s}$ is locally projected onto its dominant plane, where the dominant plane is defined by it local orientation θ_i, and s denotes the characteristic scale (or locality span) of the vertex of interest, which can be computed as explained in section 4.3.

A 2D patch (image) is obtained after projecting the 3D mesh onto the dominant plane, for which a traditional SIFT descriptor is computed [8].

4.3 Scale-Invariant 3D local descriptors

One of the approaches for selecting a subset of vertices, proposes to detect a set of so-called points of interest [8, 9], where a vertex is considered as a point of interest if it maximizes the local variation of a function of the image of 3D model, both in location and scale spaces. In the case of 3D models, this is the straightforward extension of the Gaussian scale space used by Lowe to select points of interest in 2D images [8], where the Gaussian scale space is constructed by cascading a Gaussian filter and a smoothing of the image, which create a set of filtered images called octaves.

However, an important difference between applying Gaussian filters to 2D images and 3D meshes is that points on a mesh are not necessarily uniformly spaced in the depth axis[9]. This characteristic, might harm the repeatability rate and the scale-invariant capabilities of the selected points [9, 30]. To cope with this issue, several approaches have been proposed [9, 27, 31, 32], among which, the use of local filters with uniform weights to build the mesh octaves outperforms the others methods in terms of repeatability rate and scale invariance [9].

In this approach [9], the smoothing step required to build the Gaussian scale space is given by the estimation of each vertex $\boldsymbol{v_i^{s+1}}$ at scale $s+1$ as,

$$\boldsymbol{v_i^{s+1}} = \frac{1}{|V_i^s|} \sum_{\boldsymbol{v_j^s} \in V_i^s} \boldsymbol{v_j^s}, \tag{3}$$

where, V_i^s is the set of first order neighbors of $\boldsymbol{v_i^s}$, $|\cdot|$ denotes the cardinality operator, and the summation of the vertices is performed component-wise on the x, y, and z components of $\boldsymbol{v}$.

Once the smoothing of the vertex is performed, the difference of Gaussians (DoG) is computed by,

$$d_i^s = \frac{1}{\left(s \cdot \sigma_{i,0}^2\right)} \left(\boldsymbol{v_i^s} - \boldsymbol{v_i^{s+1}}\right), \tag{4}$$

where, σ_0 is the initial variance of the integration parameter [8] and it is independently estimated for each vertex as,

$$\sigma_{i,0} = \frac{1}{|V_i^s|} \sum_{\boldsymbol{v_j^s} \in V_i^s} abs\left(\boldsymbol{v_i^s} - \boldsymbol{v_j^s}\right), \tag{5}$$

where, $abs\,(\cdot)$ denotes the absolute value operator.

Using this formulation, it is possible to select a subset of points of interest for which a local descriptor (SIFT or spin image) will be computed, i.e., those points whose d_i^s is maximal both in location and scale spaces. Section 6 will show that this scheme to detect points of interest and select their characteristic scale achieves good retrieval and classification performance when it is combined with Spin images and SIFT descriptors.

5 Experiments

This section introduces the dataset generated by our project during its initial phase, as well as the set of experiments that we have performed.

5.1 Data

Our data consists of 149 3D surfaces that were manually scanned from 16 sherds of Aztec ceramics, and annotated under different taxonomic criteria by experts in archaeology. More precisely, surfaces obtained from different views were scanned from the internal and external sides of the sherds, and the following taxonomic criteria were used to annotate them according to their visual information,

- *type*: defines the type of ceramic where the sherd comes from. At this stage of the project, three types of ceramics were available to us: Apaxtle, Cajete, and Cajete with handler (see Fig. 1).
- *side*: the side of the sherd with two possible values: internal and external.
- *form*: this label indicates the form of the ceramic where the sherds comes from, with two possible values: open and close.
- *wall*: corresponds to the type of wall of the original ceramic, with three possible values: simple wall, divergent, and straight.
- *curvature*: this label indicates whether the sherd corresponds to a flat or curved section of the ceramic. It has two possible values: NONE and YES.
- *border*: gives information about the type of border of the ceramic. It has three possible values: NONE when no border is visible in the 3D surface, straight, and round.
- *base*: this label indicates whether the sherd contains part of the base of the ceramic, and it has two possible values: NONE and YES.

Fig. 4a shows the amount of 3D surfaces scanned from each of the 16 sherds. Note that the dataset tends to be balanced in the amount of internal and external views that were scanned, despite the fact that only external views were available for sherd 11. Also, Fig. 4b shows the distribution of labels on the 149 3D surfaces. Note that only 139 3D surfaces are used for the label ‘border’, because 10 models have borders different from the three possible values (NONE, straight, round), but they have only one or two instances, which makes difficult to define relevant sets, and therefore, these instances were excluded for this taxonomy.

Although the available data after scanning the sherds contains both geometry (mesh defined by points and faces) and texture (surface photos), we started

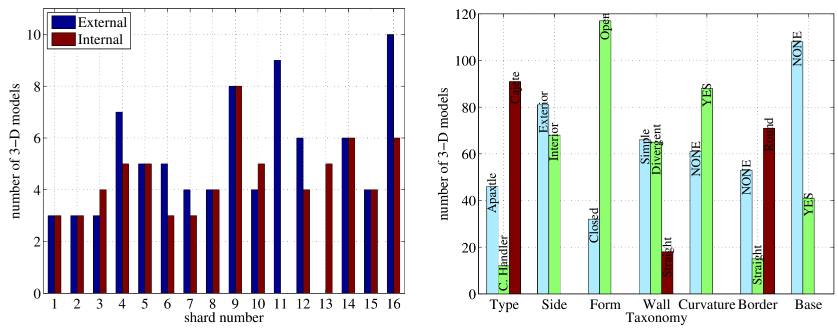

(a) Amount of 3D surfaces per sherd. (b) Amount of 3D surfaces per label.

Fig. 4: (a) Amount of 3D surfaces scanned per sherd: blue bars correspond to external views, whereas red bars indicate internal views. (b) Distribution of labels assigned to each 3D surface.

our analysis relying solely on the geometric information. For this purpose, the SIFT descriptor was tailored to work with orientation information instead of intensities, the same way as the work of Darom and Keller [9].

5.2 Experimental protocol

The experimental protocol that we followed consists of:

1. Compute the sets of local descriptors (SIFT [8] or spin images [10]) for each of the 149 3D surfaces.
2. Estimate a visual vocabulary using a subset of 700 local descriptors for each of the 16 sherds, thus 11200 local descriptors. We computed visual vocabularies of different sizes, i.e., 100, 250, 500, 1000, and 2000 visual words.
3. Build a bag-of-words representation (BoW) of each 3D surface.
4. Perform retrieval experiments by sorting the 3D surfaces in terms of visual similarity, where such similarity is computed using the BoW. Note that these retrieval experiments were performed using a leave-one-out full-cross-validation approach. We evaluated the retrieval performance using the different taxonomic criteria to define the relevant set for each 3D query surface.
5. Using PCA [33], we found the most informative principal components of the BoW's, and repeated the retrieval experiments using only them.
6. Finally, we performed a series of classification experiments using the kNN approach for the seven taxonomic criteria.

This protocol was repeated independently for each of the local descriptors and each visual vocabulary. We used a Matlab implementation [9] to compute both SIFT and spin images local descriptors. In section 6, we report our results in terms of mean average precision computed for the first 25 3D surfaces retrieved (mAP@25) and the standard recall vs average precision curves . We also report the average classification rate and show the most interesting confusion matrices.

6 Results

Table 1 shows the $mAP@25$ obtained with both the SIFT [8] and the spin images [10] local descriptors, and using different vocabulary sizes and the different taxonomic criteria described in section 5.1.

The first observation available from Table 1, is that small vocabularies are adequate for description of the 3D surfaces of potsherds for most of the taxonomic criteria, i.e., 100 and 200 visual words in most cases. Two exceptions to this observation were found using the Spin images along with the Curvature and Border information, where respectively, larger vocabularies of 2000 and 1000 visual words achieved the best retrieval performance.

Table 1: $mAP@25$ obtained with the SIFT and spin images descriptors using different vocabulary sizes and taxonomies. Best result per taxonomy is in bold.

Vocabulary	Taxonomy						
SIFT	Type	Side	Form	Wall	Curvature	Border	Base
100	**0.688**	**0.749**	0.802	**0.651**	**0.704**	**0.643**	**0.767**
200	0.677	0.734	**0.816**	0.637	0.694	0.635	0.744
500	0.649	0.721	0.812	0.605	0.687	0.633	0.736
1000	0.618	0.709	0.807	0.592	0.687	0.640	0.727
2000	0.613	0.687	0.810	0.574	0.670	0.627	0.701
Spin	Type	Side	Form	Wall	Curvature	Border	Base
100	**0.579**	0.610	**0.753**	0.526	0.619	0.533	**0.660**
200	0.577	**0.618**	0.751	**0.532**	0.617	0.540	0.655
500	0.573	0.617	0.752	0.527	0.620	0.544	0.638
1000	0.569	0.616	0.751	0.521	0.625	**0.549**	0.634
2000	0.564	0.622	0.747	0.522	**0.626**	0.544	0.626

Another observation from Table 1 is that, good retrieval performance is obtained by using the Form information to build the relevant sets, i.e., 0.816 with the SIFT descriptors and 0.753 with the spin images, which suggests that it is relatively easy to differentiate between open and close forms. However, this result must be read with caution, since the dataset is not well balanced under such a taxonomic criterion as shown in Fig. 4b, where the set of models with closed forms is about a third of the set with open forms.

The construction of the relevant set using the taxonomic criteria Side, Curvature, and Base also provided with acceptable retrieval performance, i.e., above 0.7 using the SIFT descriptors and above 0.6 with the spin images. Note that the bars in Fig. 4b indicate that all these three taxonomies define a binary classification problem, with balanced relevant sets for the Side and Curvature cases. Thus suggesting that the use of such taxonomies can provide with good hints for automatic categorization. The remaining three taxonomic criteria Type, Wall, and Border, all divide the dataset into three classes, and all obtained the poorer retrieval performance for either of the descriptors.

Finally, note that the SIFT descriptors are more appropriate than the spin images for the statistical visual description of 3D surfaces of potsherds, as they obtained better retrieval performance in all cases, i.e., about 10% better on average, which confirms observations from previous comparisons [9, 29].

Using Principal Component Analysis (PCA) [33], we performed dimensionality reduction on the BoW's and evaluated the retrieval performance. We noticed that the retrieval precision increased in all cases, specially when large vocabularies are used, and with the SIFT descriptors. Table 2 shows the best retrieval results, which on average were obtained using 1000 visual words.

Table 2: mAP@25 obtained using 1000 words and principal component analysis. The number of principal components used with each taxonomy is also shown.

Vocabulary	Taxonomy						
SIFT	Type	Side	Form	Wall	Curvature	Border	Base
1000	0.718	0.788	0.821	0.678	0.695	0.658	0.776
Num. Comp.	5	9	5	9	20	20	10
Spin	Type	Side	Form	Wall	Curvature	Border	Base
1000	0.590	0.627	0.734	0.527	0.634	0.560	0.678
Num Comp.	9	18	9	20	5	5	6

Looking at the Type column in Table 2, we can see that by using only 5 principal components from a 1000 vocabulary, higher retrieval precision is obtained, i.e., 3% higher than using 100 words, and 10% higher than using the full 1000 words vocabulary. Due to space constraints, we only present here the results with 1000 words, as they represent the case of the larger increment in retrieval precision. However, this behavior resulted true for most of our experiments.

Fig. 5a and Fig. 5b show, respectively for SIFT and Spin images, the average precision vs standard recall curves obtained with the visual vocabulary of 1000 words. This is, the average drop of the retrieval precision as more relevant instances are included within the vector of sorted elements. Note that the drop of the retrieval precision is smooth in all cases. Also note that these curves are consistent with the results shown in Table 1 and Table 2 for both descriptors, i.e., Form information provides the best retrieval performance (red-dotted line with diamond markers), whereas the Border and Wall taxonomies obtained the lowest retrieval performance (olive-dotted line with triangle markers and cyan-solid line with asterisk markers respectively).

Finally, we performed a series of classification experiments using the kNN (k=1). For these experiments, we also used the principal components of the 1000 words vocabularies, as indicated in Table 2.

Table 3 shows the average classification rate achieved by both the SIFT descriptor and the spin images. Note that, similarly to the retrieval results, classify-

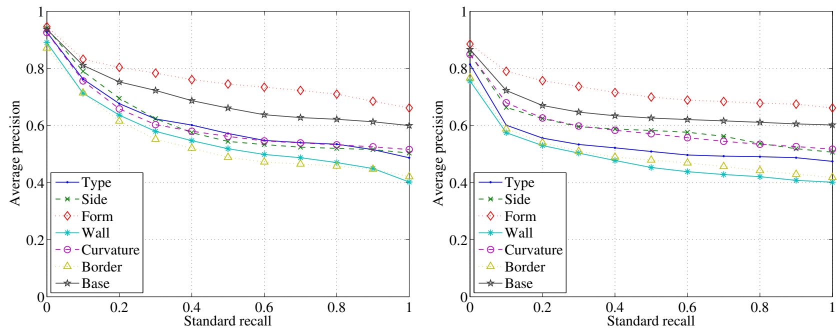

(a) SIFT descriptors. (b) Spin images.

Fig. 5: Average precision vs standard recall obtained with the different taxonomic criteria. These result were obtained using 1000 visual words.

Table 3: Average classification rate per taxonomic criteria. These results were obtained using the principal components of the 1000 words vocabularies, as indicated in Table 2.

	Type	Side	Form	Wall	Curvature	Border	Base
SIFT	0.70±0.1	0.88±0.0	0.77±0.2	0.67±0.3	0.78±0.1	0.64±0.3	0.74±0.2
spin	0.43±0.2	0.58±0.0	0.57±0.4	0.42±0.2	0.58±0.1	0.44±0.1	0.56±0.3

ing potsherds using the Side, Form, and Curvature information gives the highest accuracy. However, using the Type information as labels also gives acceptable classification rates. Also note that, the use of he Wall and Border information gives the poorest results, which contain high levels of standard deviation.

A more detailed explanation of the classification performance is shown in the confusion matrices of Fig. 6. Due to space constraints, we only show the most interesting confusion matrices.

7 Conclusions and future work

We presented a novel project towards the integration of computer vision techniques in the analysis of potsherds from the ancient Teotihuacan and Aztec cultures. Our interdisciplinary approach is rich and genuine, as it addresses needs and open problems in archaeology, i.e., those of helping the efficient analysis and categorization of potsherds.

Different from classical works on analysis of 3D potsherds, ours approaches the categorization problem under 3 assumptions: we use rotation and scale invariant local descriptors, which avoid the need of normalizing the 3D sherds under such transformations; we avoid the need of manual registration of different views of the sherds by relying on 3D surfaces; and instead of performing

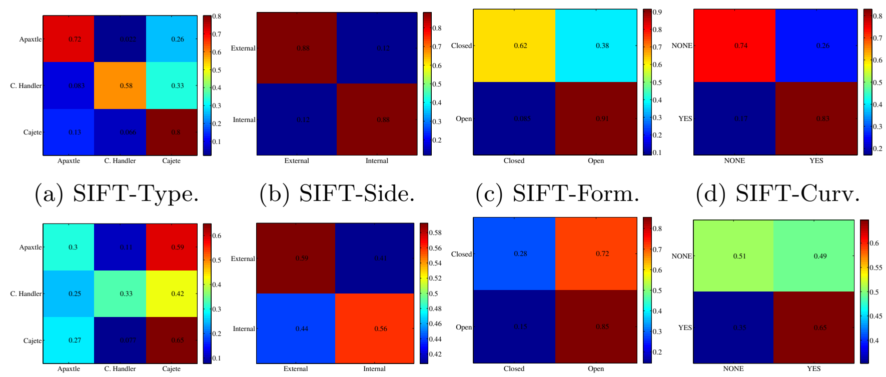

(a) SIFT-Type. (b) SIFT-Side. (c) SIFT-Form. (d) SIFT-Curv.

(e) Spin-Type. (f) Spin-Side. (g) Spin-Form. (h) Spin-Curv.

Fig. 6: Confusion matrices for Type, Side, Form, and Curvature information, computed using PCA and the number of components indicated in Table 2. Best viewed as pdf.

point-to-point matching of the 3D models, we use statistical representations that are much more efficient.

In this work, we designed a recognition system, which after a single unsupervised training, is able to categorize 3D surfaces of potsherds under different taxonomic criteria that in archaeology are used to infer the final class of sherds.

Within the context of our project, we started the collection of a new and challenging dataset of 3D surfaces of potsherds. The initial strategy is constrained to the use of Aztec III ceramics, due to the availability of this material. However, we plan to scan a much larger collection of sherds from the Teotihuacan culture, which ranks in the order of tens of thousands, and that is well documented. When this dataset is completely scanned, it will be also annotated using the different taxonomic criteria.

We presented a study of the descriptive potential of two state-of-the-art 3D descriptors for the pieces of our datasets, and combined them with the efficient bag-of-visual-words representation, and dimensionality reduction techniques based on PCA. Overall, our methodology is promising, as the results show good categorization performance using the different taxonomic criteria. More precisely, achieving an average accuracy rate of 70%, while leaving room to conduct more research in the topic. For instances, collecting more data and designing better descriptors.

Acknowledgement. This work was funded by the Swiss NSF through the project Tepalcatl-P2ELP2-152166, and supported by the project Urban Archaeology of the National Anthropology and History Institute of Mexico (INAH).

References

1. Scopigno, R., Callieri, M., Cignoni, P., Corsini, M., Dellepiane, M., Ponchio, F., Ranzuglia, G.: 3D Models for Cultural Heritage: Beyond Plain Visualization. Computer **44** (2011) 48–55
2. Kampel, M., Sablatnig, R., Costa, E.: Classification of Archaeological Fragments using Profile Primitives. In: Proceedings of Computer Vision, Computer Graphics and Photogrammetry – a Common Viewpoint, Workshop of the Austrian Association for Pattern Recognition. (2001)
3. Karasik, A., Smilanski, U.: 3D Scanning Technology as a Standard Archaeological Tool for Pottery Analysis: Practice and Theory. Journal of Archaeological Science **35** (2008) 1148–1168
4. Leymarie, F., Cooper, D., Joukowsky, M., Kimia, B., Laidlaw, D.H., Mumford, D., Vote, E.: The SHAPE Lab – New Technology and Software for Archaeologists. In: Proceedings of Computer Applications in Archaeology, BAR International Series 931. (2000)
5. Kampel, M., Sablatnig, R.: Virtual Reconstruction of broken and unbroken Pottery. In: Proceedings of the International Conference on 3-D Digital Imaging and Modeling. (2003)
6. Razdan, A., Liu, D., Bae, M., Zhu, M., Farin, G.: Using geometric modeling for archiving and searching 3d archaeological vessels. In: Proceedings of the International Conference on Imaging Science, Systems, and Technology CISST. (2001)
7. Larue, F., Di-Benedetto, M., Dellepiane, M., Scopigno, R.: From the Digitization of Cultural Artifacts to the Web Publishing of Digital 3D Collections: an Automatic Pipeline for Knowledge Sharing. Journal of Multimedia **7** (2012) 132–144
8. Lowe, D.G.: Distinctive Image Features from Scale-Invariant Keypoints. Int. J. Comput. Vision **60** (2004) 91–110
9. Darom, T., Keller, Y.: Scale-Invariant Features for 3-D Mesh Models. IEEE Transactions on Image Processing **21** (2012) 2758–2769
10. Johnson, A.E., Hebert, M.: Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes. IEEE Trans. Pattern Anal. Mach. Intell. **21** (1999) 433–449
11. Maiza, C., Gaildrat, V.: Automatic Classification of Archaeological Potsherds. In: Proceedings of the International Conference on Computer Graphics and Artificial Intelligence. (2005)
12. Quelhas, P., Monay, F., Odobez, J., Gatica-Perez, D., Tuytelaars, T.: A Thousand Words in a Scene. IEEE Transactions on Pattern Analysis and Machine Intelligence **29** (2007) 1575–1589
13. Cervantes Rosado, J., Fournier, P., Carballal, M.: La cerámica del Posclásico en la Cuenca de México. In: La producción Alfarera en el México Antiguo. Volume V of Coleccin Cientfica. Insituto Nacional de Antropología e Historia (2007)
14. Parsons, J.R.: Prehistoric Settlement Patterns in the Texcoco Region, Mexico. In: Memoirs of the Museum of Anthropology. Volume 3. University of Michigan (1971)
15. Blanton, R.E.: Prehispanic Settlement Patterns of the Ixtapalapa Region, Mexico. In: The Pennsylvania State University Occasional Papers in Anthropology. Volume 6. University Park (1972)
16. Sanders, W.T., Parsons, J., Santley, R.S.: The Basin of Mexico: Ecological Processes in the Evolution of a Civilization. New York: Academic Press (1979)
17. Parsons, J.R., Brumfiel, E.S., Parsons, M.H., Wilson, D.J.: Prehistoric Settlement Patterns in the Southern Valley of Mexico: The Chalco-Xochimilco Regions. In: Memoirs of the Museum of Anthropology. Volume 14. University of Michigan (1982)

18. Hodge, M.G., Minc, L.D.: The Spatial Patterning of Aztec Ceramics: Implications for Prehispanic Exchange Systems in the Valley of Mexico. Journal of Field Archaeology **17** (1990) 415–437
19. Hodge, M.G., Neff, H., Blackman, M.J., Minc, L.D.: Black-on-Orange Ceramic Production in the Aztec Empire's Heartland. Latin American Antiquity **4** (1993) 130–157
20. Hall, N., Laflin, S.: A Computer Aided Design Technique for Pottery Profiles. In: Computer Applications in Archaeology, Centre for Computing and Computer Science, University of Birmingham (1984)
21. Adler, K., Kampel, M., Kastler, R., Penz, M., Sablatnig, R., Schindler, K., Tosovic, S.: Computer Aided Classification of Ceramics-Achievements and Problems. In: Proceedings International Workshop on Archaeology and Computers. (2001)
22. Kampel, M., Sablatnig, R.: Rule Based System for Archaeological Pottery Classification. Pattern Recognition Letters **28** (2007) 740–747
23. Melero, F.J., Len, A.J., Contreras, F., Torres, J.C.: A new system for interactive vessel reconstruction and drawing. In: Bar International Series. Volume 1227. Oxford: Archaeopress (2004)
24. Sablatnig, R., Menard, C.: Computer based acquisition of archaeological finds: the first step towards automatic classification. In: Proceedings of the Third International Symposium on Computing and Archaeology. (1996)
25. Schurmans, U., Razdan, A., Simon, A., McCartney, P., Marzke, M., Alfen, D.V., Jones, G., Rowe, J., Farin, G., Collins, D., Zhu, M., Liu, D., Bae, M.: Advances in Geometric Modeling and Feature Extraction on Pots, Rocks and Bones for Representation and Query Via the Internet. In: Computer Applications and Quantitative Methods in Archaeology (CAA). (2001)
26. Hörr, C., Brunner, D., Brunnett, G.: Feature Extraction on Axially Symmetric Pottery for Hierarchical Classification. Computer-Aided Design & Applications **4** (2007) 375–384
27. Castellani, U., Cristani, M., Fantoni, S., Murino, V.: Sparse points matching by combining 3D mesh saliency with statistical descriptors. Computer Graphics Forum **27** (2008) 643–652
28. Gilboa, A., Karasik, A., Sharon, I., Smilansky, U.: Towards Computerized Typology and Classification of Ceramics. Journal of Archaeological Science **31** (2004) 681–694
29. Frome, A., Huber, D., Kolluri, R., Bulow, T., Malik, J.: Recognizing Objects in Range Data Using Regional Point Descriptors. In: Proceedings of the European Conference on Computer Vision (ECCV). (2004)
30. Boyer, E., Bronstein, A.M., Bronstein, M.M., Bustos, B., Darom, T., Horaud, R., Hotz, I., Keller, Y., Keustermans, J., Kovnatsky, A., Litman, R., Reininghaus, J., Sipiran, I., Smeets, D., Suetens, P., Vandermeulen, D., Zaharescu, A., Zobel, V.: SHREC 2011: Robust Feature Detection and Description Benchmark. In: Proceedings of the 4th Eurographics Conference on 3D Object Retrieval. (2011)
31. Zaharescu, A., Boyer, E., Varanasi, K., Horaud, R.: Surface feature detection and description with applications to mesh matching. In: IEEE Conference on Computer Vision and Pattern Recognition. (2009)
32. Mian, A., Bennamoun, M., Owens, R.: On the Repeatability and Quality of Keypoints for Local Feature-based 3D Object Retrieval from Cluttered Scenes. Int. J. Comput. Vision **89** (2010) 348–361
33. Jolliffe, I.T.: Principal Component Analysis. 2nd edn. Springer (2002)