Machine Vision and Applications DOI: 10.1007/s00138-009-0198-7
(The original publication is available at www.springerlink.com)

A novel framework for 3D reconstruction and analysis of ancient inscriptions

Angelos Barmpoutis[1], Eleni Bozia[2], Robert S. Wagman[2⋆]

[1] Computer and Information Science and Engineering, University of Florida, Gainesville FL 32611, USA
[2] Department of Classics, University of Florida, Gainesville FL 32611, USA

Received: January 16, 2008 / Revised version: January 29, 2009 / Accepted: April 14, 2009

Abstract Studying ancient inscriptions is based up to date mostly on observation and manual analysis by means of which epigraphists attempt to establish a geographical and chronological classification as well as to analyze the lettering techniques. In this paper we propose a novel framework for efficient 3D reconstruction of inscriptions and for statistical analysis of their reconstructed surfaces. The proposed framework employs a shape-from-shading technique to reconstruct in 3D the shape of the inscribed surfaces. The obtained surfaces are segmented into smaller box-shaped regions containing single letters. These letters are classified into groups of same characters or symbols and then an atlas (average) letter shape is created for each character. For the construction of those atlases we employ a functional minimization method that registers the surfaces of same letters to the unknown average surface, which is also estimated simultaneously. Using the estimated atlases an automated analysis of the inscribed letters is performed. This framework can be effectively used for the study of the variations of the lettering techniques within an inscription or a set of inscriptions. We applied our framework to five ancient Greek inscriptions. Our results are reported in detail and the variations found in lettering techniques are commented on by archaeologists who also validate the accuracy of our proposed method.

1 Introduction

The evolution of imaging technology has affected positively several areas in the computer vision community. During the last decades the significant progress in sensoring technology has helped us capture high-resolution multimodal images that have consequently resulted in the advancement of medical imaging (Magnetic Resonance Images [2,23]), general image and video processing [8,26,22], and virtual 3D environment synthesis (for instance, reconstruction of archaeological monuments [18, 7,21,19]).

In archaeology a crucial field of research is the study of ancient inscriptions. Some questions that archaeologists are called to answer pertain to the dating of the inscriptions, the attempt to identify the place of origin, and the analysis of the lettering techniques. So far the method that has been commonly used for the study of inscribed fragments is the following: the archaeologists use a special type of moisturized paper (squeeze) which they push on the inscribed surface using a brush specially adapted for the purpose. When the letters are shaped on the squeezed paper, the archaeologists let it dry, creating that way an impression of the inscription (see Fig. 2). The paper most in use is a type of filter paper obtained from laboratory suppliers and it may vary in thickness. A thinner paper needs less working on the stone, but may disintegrate more easily [27]. Other media have been also used in literature for this purpose, for instance latex, or liquid rubber, whose properties are discussed in [3]. By using the above technique, large collections of squeezes have been created, including squeezes of inscriptions that are now lost or destroyed, that can be found and studied in various libraries and institutes around the world.

Although squeezes are copies of inscribed surfaces, there are several problems that may rise during their use. First, their accessibility is still limited; archaeologists have to travel in order to study a variety of squeezes. Their preservation also constitutes a problem in that folding and unfolding them several times may result in their deterioration. The aforementioned problems have led us to look for a method to store and preserve squeezes more efficiently and make them also more accessible to a large number of archaeologists and epigraphists.

A simple approach that has been extensively used to capture digitally an inscribed surface is to take a picture

⋆ This research was in part funded by the Gerondelis Foundation Inc. grant to EB and the University of Florida Alumni Fellowship to AB.

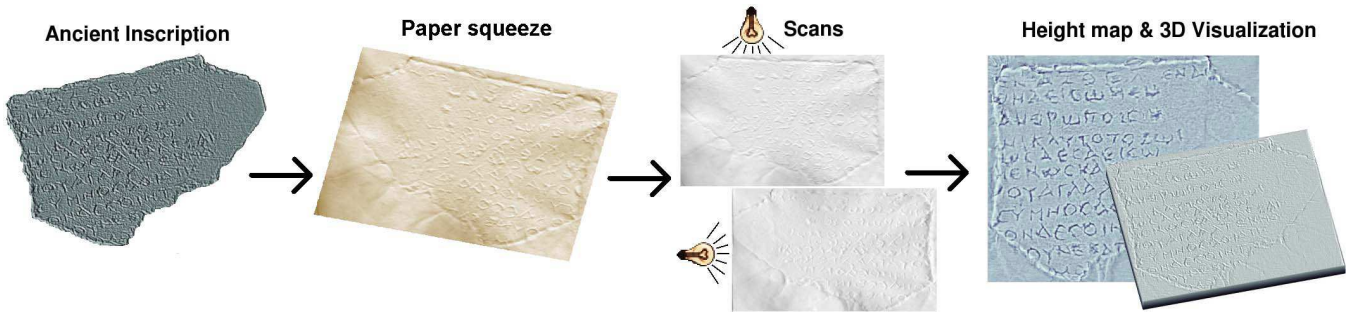

Fig. 1 Illustration of the proposed framework for 3D reconstruction of inscribed surfaces by scanning squeezes using two different lighting directions.

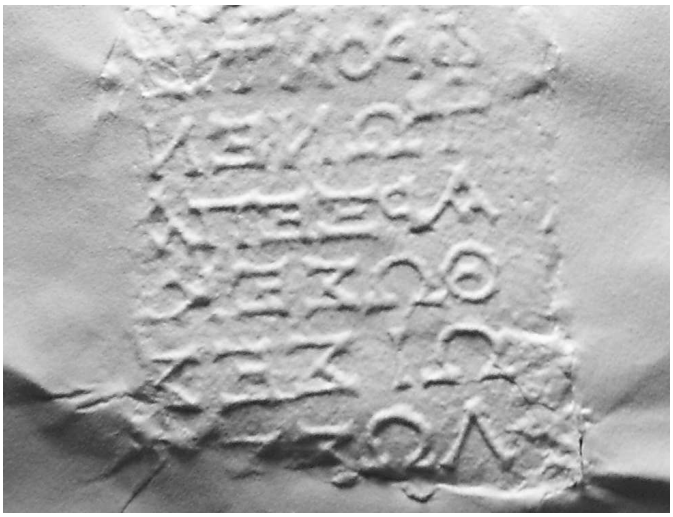

Fig. 2 An epigraphic squeeze with the impression of an inscription. (The back side is shown.)

of the surface [4,20,17,11]. Alternatively one can also scan the squeeze [24]. The problem with the aforementioned methods is that a 2D image cannot capture the 3D content of those subjects and it is also sensitive to lighting conditions and camera view point.

Several fixed pose pictures of the inscription taken under varying illumination conditions can capture the underlying local information of the inscribed surface [16]. In this technique a hemi-spherical grid of light sources illuminates the subject from a set of known lighting directions. Then the local bidirectional reflectance distribution function can be approximated by a fixed-order polynomial known as Polynomial Texture map. Although this method works nicely for surface relighting and manipulating its reflectance characteristics, the construction of the hemi-spherical grid of light sources can be expensive, considering that the size of the grid is proportional to the size of the inscription, which sometimes is several feet wide.

A video sequence or a set of 2D images can also be used for reconstructing a 3D scene [5,19]. This technique is very efficient and requires only an uncalibrated camera. Such methods have been successfully used for the 3D reconstruction of archaelogical sites, including buildings and complex statues [19]. However, the precision and detail of the obtained reconstructions is not high enough for capturing details on inscribed surfaces such as letters and symbols.

Another more sophisticated way is to use laser scanners [14,15,12], devices that can capture in a very detailed way the inscribed surface. This method has been applied to hieroglyphics and it looks promising [14]. The disadvantages, however, lie in the fact that manufacturing a laser scanner can prove costly and the transportation and use of a laser scanner is not always feasible, considering the limited accessibility of many archaeological sites.

A more cost effective method that requires only the squeezes is presented in this paper. This method requires a regular image scanner with which one can scan a squeeze from two different lighting directions. We use these images in a shape from shading technique [10] in order to reconstruct in high resolution the original 3D surface. One can store the 3D output digitally, study the squeezes, and also analyze them quantitatively using automated statistical tools. An illustration of our proposed framework is presented in Fig. 1. We have applied this method on 5 inscriptions from Epidauros in Greece and we present extensive experimental results from their 3D reconstruction and the analysis of their lettering technique.

The contributions of our proposed method are the following: it is the first framework reported in literature that converts and stores squeezes in a 3D digital form. It maximizes also the utilization of squeezes since they can be more effectively studied using different visualizations (viewing angles, lighting conditions) and they can also be copied and distributed to the research community. Finally, the automated analysis, as presented in this paper, can produce results that were very difficult, if not impossible, to get with the standard techniques.

The rest of the paper is outlined as follows. In Sec. 2 we present the method for the 3D reconstruction of ancient inscriptions. Section 3 is divided into several subsections describing the post-processing steps that we followed for the statistical analysis of the reconstructed in-

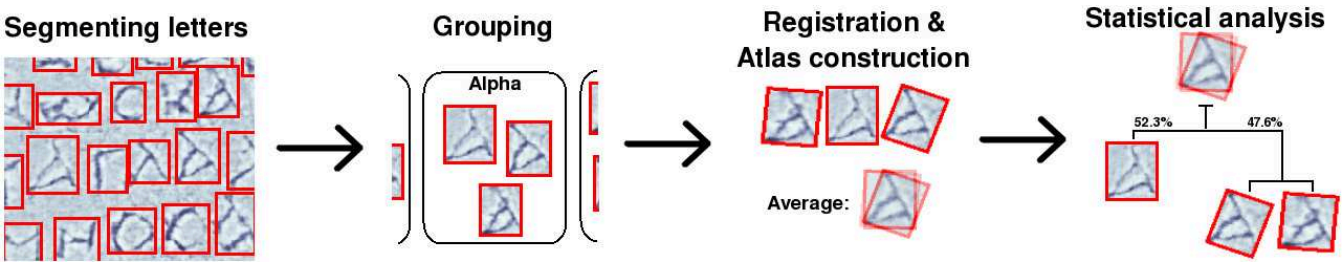

Fig. 3 Illustration of the steps followed for the statistical analysis of the lettering technique in the reconstructed inscriptions.

scriptions (Letter segmentation in 3.1, Grouping in 3.2, Registration and Atlas construction in 3.3 and Clustering in 3.4). In Sec. 4 we present our experimental results and in Sec. 5 we conclude.

2 Surface reconstruction

Surface reconstruction using photometric stereo methods has been studied extensively in computer vision literature [10]. In those methods the information of a 3D surface is recovered by using a set of fix-posed images of the subject taken under different known illumination conditions. The disadvantage of the aforementioned methods lies in the fact that the results depend on the Bidirectional Reflectance Distribution Function (BRDF) of surface's material. For instance, non-Lambertian surfaces with specularities require a large number of acquired images in order to lead to realistic results [1,22]. In our application the depicted subject is a paper squeeze (a type of filter paper), whose specular component is approximately zero and therefore its BRDF can be well represented by the diffusive parameter of the Phong reflectance model (also known as Lambertian model [1]). Thus, accurate results can be obtained by applying a shape-from-shading method (presented in this section) using only two scanned images of each squeeze, which are sufficient to reconstruct the unknown parameters in this reflectance model.

Assume that the normal vector of the underlying surface S at the (i,j) location is given by $[p(i,j)\ q(i,j)\ 1]$, where $p(i,j) = \frac{\partial S(i,j)}{\partial x}$ and $q(i,j) = \frac{\partial S(i,j)}{\partial y}$ are the x, y gradients of the surface respectively.

Each scanned image of the surface can be modeled using the Phong reflectance model. In this model the reflectance of a Lambertian surface (no specularity) is given by

$$R(i,j) = k_a l_a + k_d l_d (\mathbf{L} \cdot \mathbf{N}(i,j)) \tag{1}$$

where $\mathbf{L}$ is the 3-dimensional vector of the direction of the light beam at (i,j), $\mathbf{N}(i,j)$ is the normal vector of the surface at (i,j), l_a and l_d are the ambient and diffusion component of the light and finally k_a and k_d are the ambient and diffusion components of the surface material respectively.

Given a set of N scanned images $I_1, I_2, \ldots, I_N$ associated with light source directions $L_1, L_2, \ldots, L_N$ respectively we need to estimate the unknown surface gradients p and q by minimizing the following energy

$$E(p,q) = \sum_{n=1}^{N} \int_{i,j} ((I_n - R)\sqrt{p^2+q^2+1})^2$$

$$= \sum_{n=1}^{N} \int_{i,j} ((I_n - c_a)\sqrt{p^2+q^2+1} - c_d(pL_x^n + qL_y^n + L_z^n))^2 \tag{2}$$

where $c_a = k_a l_a$, $c_d = k_d l_d$, and L_x^n, L_y^n, and L_z^n are the x, y, z components of the direction of the n^{th} light source. Note that I_n, R, p and q are all 2-dimensional functions and in Eq. 2 are being integrated over their domain. Equation 2 is minimized when R comes close to I_n and the role of the factor $\sqrt{p^2+q^2+1}$ is to maintain numberical stability of the functional minimizing method.

In Eq. 2 we can also add a regularization term for smoothing of p and q across the lattice. The regularization term can be expressed by

$$\int_{i,j} \left(\frac{\partial p}{\partial x}\right)^2 + \left(\frac{\partial p}{\partial y}\right)^2 + \left(\frac{\partial q}{\partial x}\right)^2 + \left(\frac{\partial q}{\partial y}\right)^2 \tag{3}$$

The standard role of the regularity term (Eq. 3) is to remove part of the high frequencies from the data. In our case, these frequencies correspond to fine details of the squeezes, which more often than not are undesired noise on the surface of the paper. However, some of the useful details in the data may also be faded out due to regularization. Hence one needs to estimate carefully the desired degree of smoothness or incorporate a more sophisticated noise removal technique (i.e. anisotropic smoothing). In our experiments we acquired data which did not contain noise artifacts and therefore we did not impose any regularization constraints.

Given the surface gradients p and q, we can reconstruct the unknown surface S by minimizing the following energy

$$E(S) = \int_{i,j} \left(\frac{\partial S}{\partial x} + p\right)^2 + \left(\frac{\partial S}{\partial y} + q\right)^2 \tag{4}$$

3 Statistical analysis

In the previous section we estimated the unknown surface gradient fields p and q from a set of scanned images, and then we reconstructed the unknown surface itself. The reconstructed surface can be further processed in order to analyze the variability of the inscribed letters within the same inscription or in a set of inscriptions. In order to do such a statistical analysis a sequence of preprocessing steps is required to be performed. These steps involve segmentation of the letters within an inscription (sec. 3.1), grouping of the segmented surfaces representing the same character or symbol (sec. 3.2), registration of the letter-surfaces belonging to the same group and finally construction of an atlas (average surface) for each type of character (3.3). The procedures followed in those steps and then the statistical analysis performed after that (sec. 3.4) are described in this section.

3.1 Letter segmentation

In section 2, we reconstructed the 3D model of an inscribed surface from multiple scans of the corresponding squeezed paper. In this step we need to segment the obtained surface into small rectangular regions containing each inscribed letter.

To simplify the procedure we assume that the reconstructed inscription has the appropriate x-y orientation, such that each line of letters is parallel to x-axis. By using this assumption each letter can be contained into a rectangular region oriented parallel to the Cartesian axes.

Each rectangular region can be parameterized using 4 real numbers x_1, x_2, y_1 and y_2 denoting the left most and right most x-coordinate of the region and the upper and lower y-coordinate of the segmented region respectively. By initializing those parameters in such a way that the resulting region is of a predefined size and centered in a specific location, we can variate slightly the rectangular region in order to minimize the quantity

$$\int_{x_1}^{x_2} p(x,y_1)^2 + p(x,y_2)^2 + \int_{y_1}^{y_2} q(x_1,y)^2 + q(x_2,y)^2 \quad (5)$$

By minimizing the above functional we minimize the surface gradient along the boundaries of each rectangular box. We can also add more terms in the above functional in order to force the segmented regions to be of a certain size.

The initialization of the first rectangular segment is performed manually by marking the upper left character on the inscription, using a graphical user interface. After having estimated the parameters of the first rectangular segment, we can automatically initialize the rectangular box containing the next letter (on the right) by centering it to $((x_2 - x_1)/2 + c, (y_1 + y_2)/2)$, where c is a predefined constant. A similar propagation rule can be defined to initialize the letter on the next line of letters. In our experiments the segmentation process was propagated automatically and the only user input was the initialization of the algorithm at the location of the upper left character.

3.2 Letter Grouping

After having segmented all the inscribed characters or symbols we end up with a set of smaller surfaces parameterized also as gradient fields p and q. In this step a label must be assigned to each of these surfaces indicating a group of characters. All the regions representing the same inscribed character must be assigned to the same label. Initially we performed this proccess manually, by using an appropriately designed user interface. After having constructed the groups of the same characters we computed the average letter of each group by using the procedure discussed in section 3.3. After this step we used the estimated average letters as prior models in order to perform the grouping step automatically. This automatic proccess is described below.

After the construction of each rectangular segment using the method described in section 3.1, the corresponding segmented surface is registered using affine transformation A_n to each of the average surfaces $(\hat{p}_n, \hat{q}_n)$ representing the prior models $n = 1, \ldots, N$. Then we compare the average models with the registered surfaces by computing the following weights

$$w_n = \int_{i,j} (\hat{p}_n - p \circ A_n)^2 + (\hat{q}_n - q \circ A_n)^2 \quad (6)$$

where (p, q) are the gradient fields of the new segmented region, $(\hat{p}_n, \hat{q}_n)$ are the prior model gradient fields and A_n is the estimated affine transformation to each model.

Finally, we assign to the segmented region the label which corresponds to the smaller weight w_n.

$$label = argmin\,(w_1, w_2, \ldots, w_N) \quad (7)$$

More details regarding the affine transformation, registration and atlas construction are given in the following section.

3.3 Registration and atlas construction

In this step, given a set of N similar surfaces parameterized as gradient fields p and q, we need to estimate the unknown average surface $\hat{p}, \hat{q}$ and the unknown transformations that map (transform) each surface to the estimated average. Both tasks, registration and atlas construction, can be done simultaneously following known image processing methods [13] modified to meet the needs of our particular application.

In our application we consider only rigid transformations and scale because of the nature of the inscribed

symbols. The inscribing tools do not allow flexibility of the shape of the inscribed letters like those in manuscripts. Therefore, the same characters inscribed by one person using the same lettering technique should differ from each other by (approximatelly) a rigid transformation and scale. Any affine transformation of the 2-dimensional Cartesian coordinates (x, y) can be written in the form $A(x, y) = \mathbf{M}[x\ y]^T + \mathbf{T}$ where $\mathbf{M}$ is a 2×2 transformation matrix and $\mathbf{T}$ is a 2×1 translation vector. In our case we parameterize $\mathbf{M}$ using a rotation angle θ and a scale factor s.

The coefficients of the unknown transformations can be estimated simultaneously with the unknown atlas $\hat{p}, \hat{q}$ by minimizing the following energy function

$$E(\hat{p}, \hat{q}, A_1, \ldots) = \sum_{n=1}^{N} \int_{i,j} (\hat{p} - p_n \circ A_n)^2 + (\hat{q} - q_n \circ A_n)^2 \tag{8}$$

Note that in the argument of Eq. 8 A_n denotes minimization over the two components of translation T_n, the rotation angle θ_n and the scale factor s_n.

3.4 Clustering

After having registered the surfaces representing the same inscribed symbol, we can perform any statistical analysis (e.g. Independent Component Analysis) or clustering method to explore the variations observed in the inscribed letters of the same epigraph or a set of epigraphs [6]. For this purpose we need to define a distance measure between surfaces representing the same character. The distance between (p_1, q_1) and (p_2, q_2) can be defined as a metric on the space of registered surfaces $(p_1 \circ A_1, q_1 \circ A_1)$ and $(p_2 \circ A_2, q_2 \circ A_2)$. In this case the distance between the registered surfaces is given by

$$\int (p_1 \circ A_1 - p_2 \circ A_2)^2 + (q_1 \circ A_1 - q_2 \circ A_2)^2 \tag{9}$$

where A_x denotes the estimated affine transformation (sec. 3.3) and the integration is over the 2D spatial domain.

By using the metric defined above we can compute the $N \times N$ distance matrix $\mathbf{D}$ whose elements are $D_{i,j} = dist(\{p_i, q_i\}, \{p_j, q_j\})$, where $\{p_1, q_1\}, \ldots, \{p_N, q_N\}$ is the set of N surfaces belonging to the same group, i.e. they represent the same inscribed character.

The obtained distance matrix $\mathbf{D}$ can be used by any unsupervised clustering method that clusters the N surfaces into categories by grouping together those with similar characteristics. In our experiments we implemented the method of Agglomerative Hierarchical Clustering [9] and we briefly review it here for convenience. This method classifies N elements into a predefined number of c classes $(c < N)$, by using the distance matrix $\mathbf{D}$. Each element is initialized as a separate cluster, and in each iteration the two closest clusters merge (therefore the algorithm runs for $N - c$ iterations). There are several merging criteria that can be used in the clustering process, and their role is to specify the distance between two clusters. In Sec. 4 we present five different clustering rules and we employ them in the clustering algorithm.

Fig. 4 Example of 3D reconstruction results. Left: The two scanned images of the squeeze used in our method. Right: The height-map obtained after the reconstruction (Dark intensities represent deeper inscribed locations.)

In our experiments we used $c = 1$ in order to be able to produce a dendogram of height $h = N - 1$, where in the root there is a cluster containing all the elements, in the next level this cluster is splitted into two children (smaller clusters) and finally we end up with the last level which contains N leafs (single element clusters). At any level the distance between nearest clusters can provide the dissimilarity value for that level. Such a dendogram is a useful tool for exploring the variability of the inscribed letters, since one can observe the dissimilarities between clusters and finally come to a conclusion regarding what is the number of significantly different clusters (here group of letters).

4 Experimental Results

In this section we present experimental results of our framework using real datasets from ancient inscriptions. The epigraphical material used for this experiment consists of five inscribed fragments from the archaeological site of Epidauros in Southern Greece, containing religious hymns for Asclepius and other deities (IG IV I 2, 129-135; SEG 30, 390 in [25]). Although found in different locations, these inscriptions show close affinities in content (sacred poetry), material (red local limestone), and writing (uncial lettering of the third century CE). The analysis (presented in this section) of the letterforms captured from the paper squeezes supports the idea that all of them came from the same monument, probably an inscribed wall from one of the buildings in the sanctuary's main area.

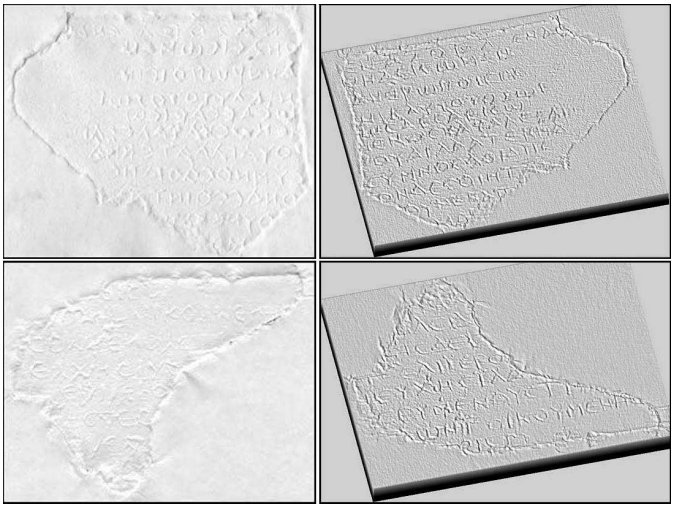

Fig. 5 3D visualization of two of the reconstructed inscribed surfaces (right). One of the two scanned images used by our method is shown on the left plates.

4.1 Surface reconstruction

In our experiments we used a standard picture scanner (EPSON Perfection 1240) set to grayscale scanning. Each squeeze was scanned twice by placing it on the scanner first horizontally and then vertically. This way the two acquired images depict the squeeze lighted from the top and from the side respectively. In each case the lighting direction was $[0\ sin(\pi/9)\ cos(\pi/9)]$ and $[sin(\pi/9)\ 0\ cos(\pi/9)]$ respectively and was determined geometrically by measuring the distance between the light source and the sensors and the distance between the sensors and the scanning surface.

The size of each image was about 1000×1000 px. Each pair of images was automatically registered since the subject was fixed in the upper left corner of the scanner and rotated by 90^o in the second image; therefore the transformation was known. Figure 4 shows the two images acquired by scanning one of the squeezes (left plate). By observing the two images one can see that in the upper image the light source is on the left of the squeeze while in the lower image the light source is located on its upper side. By applying the method presented in Sec. 2 we obtained the 3D surface of the squeeze shown as a height-map in Fig. 4 right. By inspecting visually the result, we notice that our method estimated in high detail the anaglyph of the inscribed surface, showing in darker intensities the deeper inscribed regions of the inscription. Note that in the scanned images, which is one of the common ways of publishing inscriptions [24], the inscribed text is not as legible as in the height-map obtained by our framework.

Figure 5 (right) shows 3D visualizations of two of the reconstructed surfaces. Rendering the inscriptions with different virtual illuminations and viewing angles makes the use of a squeeze more effective and allows the archaeologists to share digital copies of the squeezes without losing any information. Thus by using our proposed framework the digital libraries of scanned squeezes (regular 2D images) which are commonly used by archaeology scholars can easily be replaced by databases of 3D squeezes, without the need of any additional equipment.

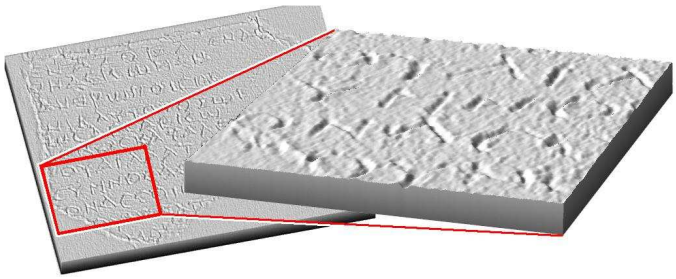

Fig. 6 Detail of a reconstructed inscription shown enlarged in 3D.

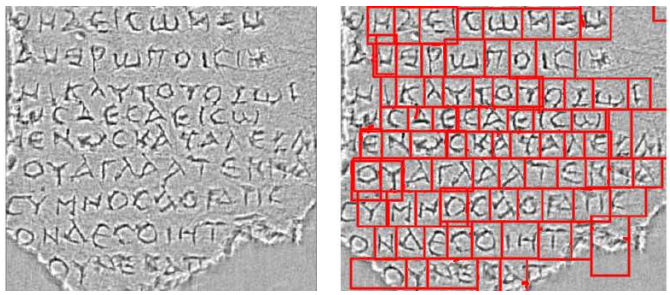

Fig. 7 Letter segmentation example. Left: The height-map of an inscription estimated by our method. Right: The obtained letter segmentation result.

Finally, figure 6 presents a zoomed region of a reconstructed inscription showing details of the inscribed surface.

4.2 Statistical analysis

After having estimated the 3D surfaces of the squeezes, we applied the method presented in Sec. 3.1 for segmenting the characters or symbols contained in each inscription. Figure 7 shows an example of letter segmentation in one of our datasets. On the left plate of the figure the estimated height-map is presented, and on the right plate the segmentation result is shown by displaying red boxes around each segmented character.

Then, the letters in our datasets were grouped by following the method presented in Sec. 3.2. Figure 8 shows four sets containing the characters 'alpha', 'epsilon', 'sigma', and 'upsilon'. Moreover, in order to perform a statistical analysis of the lettering technique, the elements of each set were registered to each other by using the method presented in Sec. 3.3. The average letters estimated for each group of characters are presented in the lower row of Fig. 8 as height-maps.

By observing the height-maps of the average letters one can make various remarks regarding the variability of the lettering schemes. For instance, by inspecting vi-

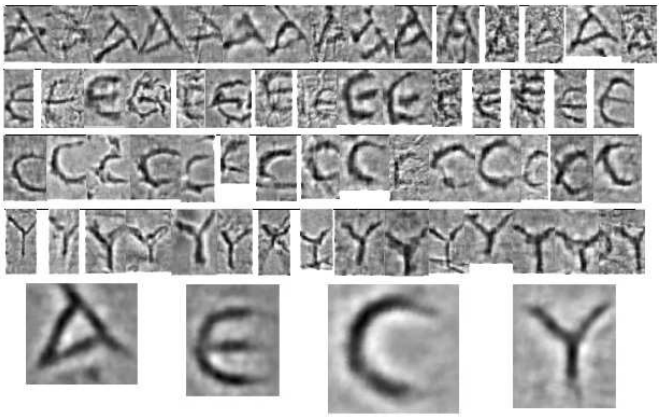

Fig. 8 The four upper rows show the elements of 4 different groups of characters extracted from our datasets. The last row shows the atlases (average element) estimated for each group.

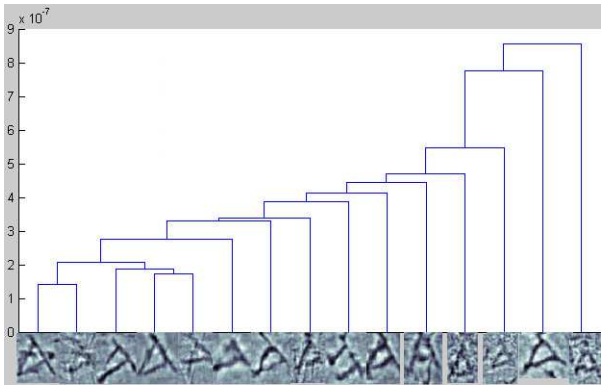

Fig. 9 Dendrogram showing the affinities between the elements of the group of character 'alpha'.

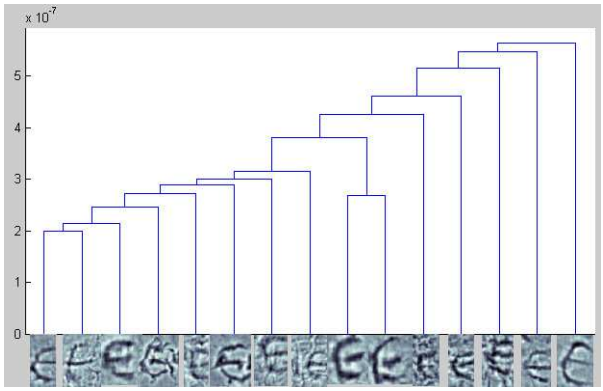

Fig. 10 Dendrogram showing the affinities between the elements of the group of character 'epsilon'.

sually the atlas of the character 'epsilon' (Fig. 8) one can observe the existence of some 'epsilons' that have the intermediate line detached from the rest of the letter structure. For this reason the height-map shows more bright intensities in the region where the intermediate line of 'epsilon' touches the vertical one. Furthermore, some of the 'epsilons' appear more elongated than the others and this is why the atlas shows extensions of the horizontal lines on the right side. Similar remarks can also be made by observing the rest of the estimated atlases. Such results were obtained up to date by epigraphists after manual observation of each individual letter in each inscription, which is a time consuming process. The automated analysis performed by our proposed framework increases significantly the speed of this process and it is also a tool that assists the scholars in the evaluation of their results.

In addition to the visual observations, our method can also produce quantitative results, as presented in Sec. 3.4, regarding the variability of the lettering technique. This step produces a dendrogram showing the affinities of the individual elements of each group of characters. In our experiments we used the nearest distance between the elements of two clusters as the merging criterion in this process. Figures 9 and 10 show the dendrograms produced for the letters 'alpha' and 'epsilon' respectively. In the plotted dendrograms one can observe the formation of smaller subgroups of characters that have a special characteristic in their structure. For instance the existence of some 'epsilons' (discussed earlier) whose intermediate line is detached from the rest of the letter can also be observed in the dendrogram of Fig. 10. In the center of this plot we can see that two letters with the above structure were automatically grouped together. Despite this and some other minor subgroups found in our experimental results, there were no other larger subgroups formed which supports the argument that the lettering technique used in all the processed inscriptions were similar.

In order to test the clustering performance, we repeated our experiments using the following five clustering criteria. We merged together those two clusters that have the smallest a) nearest distance between their elements, b) furthest distance between their elements, c) average distance between their elements, d) distance between their centers and e) within-cluster sum of squares of distances. Furthermore in order to test how many clusters the data contain, we repeated the experiments using different predefined values of the expected number of clusters and each time we computed the percentage of the variance explained by the clusters against the total variance in the dataset. The results are depicted in Fig. 11. All methods gave similar clustering results, which demonstrate robustness of the clustering step with respect to the chosen merging criterion. We should also state that the corresponding computed dendrograms were similar to those shown in Fig. 9 and 10.

We validated our framework by comparing the quantitative results shown in Fig. 11 with the expected outcome based on archaeological evidence. These evidence indicate that all the inscriptions belong to the same monument and were probably manufactured by the same workshop [25]. Therefore we expect the observed vari-

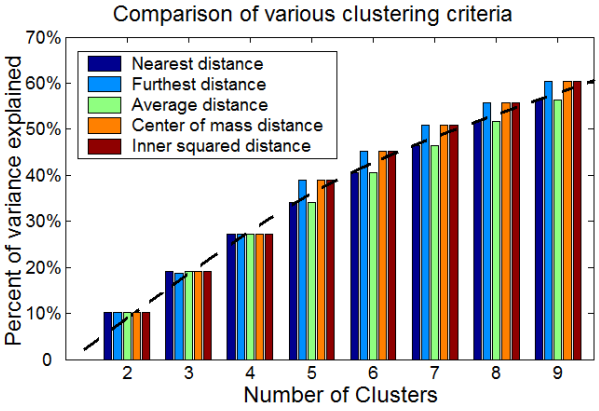

Fig. 11 Comparison of various clustering criteria. The y-axis shows the percentage of the variance explained by the clusters against the total variance in the dataset using different predefined values of the expected number of clusters (shown in x-axis). The dashed line was obtained by clustering synthetic data from a single multivariate isotropic Gaussian distribution.

ation on the letterforms to be relatively small (i.e. all letters belong to a single cluster). This observation is consistent with the obtained visual and quantitative results. First, by observing the dendrograms we notice that there are no significant formations of subclusters, which indicates that all the elements belong to the same group. This result is also validated quantitatively. We generated multi-dimensional data from a multi-variate isotropic Gaussian distribution and we clustered them using the Agglomerative hierarchical clustering. Regarding the construction of the synthetic samples, we should note that the mean and standard deviation of an isotropic Gaussian distribution does not affect the clustering performance. The percentage of the variance explained by the clusters against the total variance in the synthetic dataset is plotted in Fig. 12 (dashed line). By observing this figure, we can see that the results from the synthetic data agree with the results from the real epigraphic data, which also indicates that there are no different types of letterforms in the data.

Finally, Table 1 presents the time complexities for each individual step of the framework. According to this table the total running time in our experiments was about 16 min. For different experimental setups the reader can get an idea of the running time by observing the corresponding complexities.

5 Conclusions

In this paper we presented a novel framework for 3D reconstruction and analysis of ancient inscriptions. Our framework consists of several steps involving shape reconstruction from shading, letter segmentation, and clustering. In this framework we efficiently reconstruct in 3D inscribed surfaces by using two scanned images of the squeezed papers. There is no need for using complicated 3D scanning devices in the archaeological sites, which some times may not even be feasible. The advantages of 3D reconstruction in epigraphical research are numerous. It provides a way to store digitally the squeezes, eliminating the possibility of any damage, and also to read more clearly the inscribed text, especially in the case of badly weathered or damaged epigraphs. Moreover, the analysis and study of the lettering schemes and their variability can be performed very efficiently by using the proposed framework, which is another contribution of our method. A specialist, when using our stastistical tool, can come easily to accurate scientific conclusions, avoiding the time-consuming manual analysis and study that have been commonly used up to date.

We applied our framework to five ancient Greek inscriptions. Our results were reported in detail and the variations found in lettering techniques were commented on by archaeologists who also validated the accuracy of our proposed method.

In our future work, we will employ our framework to construct a library of 3D epigraphs. The obtained library of datasets will be used for the recovering of missing fragments from damaged inscriptions and also for the automatic dating of each inscription.

Table 1 Complexities and time performances. S is the size (resolution) of the images, N is the number of the images, K is the number of iterations at each minimization algorithm, and M is the number of inscribed characters contained in the images.

Algorithm	Complexity	Time
Acquisition	$N \times S$	≈ 40 sec
3D reconstruction	$K \times N \times S$	≈ 4 min
Segmentation	$K \times N \times S$	≈ 60 sec
Registration	$K \times M^2 \times S^2$	≈ 10 min
Clustering	M^2	≈ 1 sec

References

1. R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, 2003.
2. P. J. Basser, J. Mattiello, and D. Lebihan. Estimation of the Effective Self-Diffusion Tensor from the NMR Spin Echo. *J. Magn. Reson. B*, 103:247–254, 1994.
3. C. W. Beck. Synthetic elastomers in epigraphy. *Archaeological Institute of America*, 67:413–416, 1963.
4. E. Bozia, G. Sangco, and R. S. Wagman. A new dedication by Diogenes and other unpublished inscriptions from Epidauros. *Zeitschrift für Papyrologie und Epigraphik*, 160:120–122, 2007.
5. K. Cornelis, M. Pollefeys, M. Vergauwen, and L. J. Van Gool. Augmented reality using uncalibrated video sequences. In *SMILE*, pages 144–160, 2000.
6. R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley.

7. A. Ercil. Solving the 2D puzzle problem for automated reconstruction of archeological findings. In *Proceedings of SIU*, pages 164–167, 2003.
8. S. C. Foo. A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation. *Master's thesis, Program of Computer Graphics, Cornell University, Ithaca, NY*, 1997.
9. R. J. Gil-Garcia, J. M. Badia-Contelles, and A. Pons-Porrata. A general framework for agglomerative hierarchical clustering algorithms. *Pattern Recognition, 2006. ICPR 2006. 18th International Conference on*, 2:569–572, 2006.
10. B. K. P. Horn. *Robot Vision*. McGraw Hill: New York.
11. D. D. Hughes. A heroizing gravestone from Thera. *Zeitschrift für Papyrologie und Epigraphik*, 133:122, 2000.
12. K. Ikeuchi, K. Nishino, and A. Nakazawa. Towards the digital archive of cultural heritagepreservation and restoration of ancestral assets through observation. In *The 8th International Conference onVirtual Systems and Multimedia*, 2002.
13. S. Joshi, B. Davis, M. Jomier, and G. Gerig. Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage*, 23(1):151–160, 2004.
14. G. V. Landon and W. B. Seales. Petroglyph digitization: enabling cultural heritage scholarship. *Machine Vision and Applications*, 17:361–371, 2006.
15. M. Levoy et al. The digital Michelangelo project: 3D scanning of large statues. In *Procedings SIGGRAPH*, 2000.
16. T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *Computer Graphics, Proceedings of ACM Siggraph*, 2001.
17. G. Mennella. Una nuova dedica a Maioriano e un probabile corrector Lucaniae et Brittii nel 459. *Zeitschrift für Papyrologie und Epigraphik*, 133:237–242, 2000.
18. M. Pollefeys, L. J. Van Gool, I. Akkermans, D. De Becker, and K. Demuynck. A guided tour to virtual Sagalassos. In *VAST '01: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 213–218, 2001.
19. M. Pollefeys, L. J. Van Gool, M. Vergauwen, K. Cornelis, F. Verbiest, and J. Tops. Image-based 3D acquisition of archaeological heritage and applications. In *VAST '01: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, pages 255–262, 2001.
20. W. K. Pritchett. Photographs of the Choiseul marble. *Zeitschrift für Papyrologie und Epigraphik*, 41:149, 1981.
21. W. Shao and D. Terzopoulos. Populating reconstructed archaeological sites with autonomous virtual humans. In *Intelligent Virtual Agents*, pages 420–433, 2006.
22. A. Shashua. On photometric issues in 3D visual recognition from a single 2D image. *IJCV*, 21(1-2):99–122, 1997.
23. E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *J. Chem. Phys.*, 42 (1):288–292, 1965.
24. E. Vanderpool and W. P. Wallace. The sixth century laws from eretria. *Hesperia*, 33(4):381–391, 1964.
25. R. S. Wagman. *Inni di Epidauro*. Biblioteca di Studi Antichi, Pisa, 1995.
26. H. Wang, S. Li, Y. Wang, and J. Zhang. Self quotient image for face recognition. In *ICIP*, pages II: 1397–1400, 2004.
27. A. G. Woodhead. *The Study of Greek Inscriptions*. Cambridge University Press, 1967.