

# Petroglyph Classification using the Image Distortion Model

Vincenzo Deufemia<sup>1</sup>, Luca Paolino<sup>1</sup> and Henry de Lumley<sup>2</sup>

<sup>1</sup>University of Salerno, Fisciano(SA), Italy

<sup>2</sup>Laboratoire Départemental de préhistoire du Lazaret, Nice, France

---

## Abstract

*Petroglyphs are prehistoric engravings in stone unrevealing stories of ancient life and describing a conception of the world transmitted till today. The great number of sites and the high variability in the artifacts makes their study a very complex task. Thus, the development of tools which automate the recognition of petroglyphs is essential not only for supporting archaeologist to understand petroglyph symbols and relationships, but also for the anthropologists who are interested in the evolution of human beings. However, many challenges exist in the recognition of petroglyph reliefs mainly due to their high level of distortion and variability. To address these challenges, in this paper we present an automatic image-based petroglyph recognizer that focuses on the visual appearance of the petroglyph in order to assess the similarity of petroglyph reliefs. The proposed matching algorithm is based on an image deformation model that is computationally efficient and robust to local distortions. The classification system has been applied to an image database containing 17 classes of petroglyph symbols from Mount Bego rock art site achieving a classification rate of 68%.*

Categories and Subject Descriptors (according to ACM CCS): I.4.9 [Image Processing and Computer Vision]: Applications—I.5.1 [Pattern Recognition]: Models—Structural

---


## 1. Introduction

Petroglyphs are prehistoric engravings in stone unrevealing stories of ancient life and describing a conception of the world transmitted till today. Although they may seem as durable as the rock they reside on, petroglyphs are inevitably deteriorated by natural causes as well as vandalism. As an example some of them have been destroyed by tourists, while others are disappearing due to acid rain and sunshine. Hence, it is very important to preserve the petroglyphs trying to identify and archive them for future generations.

Many research challenges have to be addressed for the digital preservation of petroglyphs, including the integration of data coming from multiple sources and the correct interpretation of drawings. The IndianaMAS project is aiming to integrate heterogeneous unstructured data related to rock carvings into a single repository, organizing classified data into a Digital Library, interpreting data by finding relationships among them, and enriching them with semantic information [DPT\*12, MDM\*12]. These goals can be achieved by developing robust automated methods to classify the petroglyphs based on their shapes, and retrieve similar petroglyphs from different archives of petroglyph images.

From a pattern recognition point of view, the classification of petroglyphs from relief images represents a challenging task due to the high level of variability in the drawing process [ZWKL11]. Indeed, the engravings were carried out by cast percussion and in many cases two symbols in the same class have many differences. As an example, the two personages in Fig. 1(a) and 1(b) belong to the same class but have many differences in their shapes [dLE09]. The classification process is made even harder by the fact that petroglyphs have undergone a degradation process that makes them messy and/or incomplete. As an example, Fig. 1(c) depicts the relief of an ax with an extremely eroded handle that is very different from the standard handle reliefs as shown in Fig. 1(d).

In this paper we propose a classifier for petroglyph symbols based on a flexible image matching algorithm. The idea is to measure the similarity between petroglyph by using a distance derived from the image deformation model (IDM) [KDBG07], which has been successfully applied to handwritten character recognition [KGN04] and shown very good retrieval quality in the medical automatic annotation task at ImageCLEF 2005 [DWK\*05]. Such a distance measures the displacements of single pixels between two images



**Figure 1:** Two personages and two axes reliefs from Mont Bego site [dLE09].

within a warp range and taking into account the surrounding pixels (local context). This method is well-suited for petroglyph reliefs since it is less sensitive to local changes that often occur in the presence of symbol variability, making our method tolerant to the types of visual variations depicted in Figure 1.

The proposed classification system has been evaluated on a dataset containing 17 classes of petroglyph symbols from Mount Bego rock art site. The achieved results demonstrate that the proposed approach reaches a classification rate of 68%, which represents an improvement over a previous distance proposed for petroglyph recognition [ZWKL11] of about 33%.

The remainder of this paper is structured as follows. Background information and an overview of related work in petroglyph classification is given in Section 2. A description of the proposed petroglyph classification methodology based on IDM is given in Section 3. Experiments on a dataset of Mount Bego petroglyph reliefs are reported in Section 4. Conclusions are finally drawn in Section 5.


## 2. Related Work

### 2.1. Rock Art

Petroglyphs are a form of prehistoric art found in many cultures around the world and at many times. There are many theories to explain their purpose, depending on their location, age, and the type of image. Some petroglyphs are thought to be astronomical markers, maps, and other forms of symbolic communication, including a form of “pre-writing”. The form of petroglyphs is described by a variety of terms in the archaeological literature. One of the most common forms of rock art around the world are the *anthropomorphic* depictions. They are pictures that resemble humans, but sometimes can represent something else, such as the personification of a spirit or other nonliving thing. Other common images are animals, weapons, and tools.

In this work, we experimented our approach on the reliefs collected and catalogued from Mont Bego, in the extreme south-east of France, which due to the richness of the place in both qualitative and quantitative terms it is ideal for analysis. Archaeologists consider this place as an incredibly

valuable source of knowledge, due to the up to 40,000 figurative petroglyphs and 60,000 non-figurative petroglyphs [dLE09]. The figurative petroglyphs represent corniculates, harnesses, daggers, halberds, axes, reticulates, rectangular or oval shaded zones, and anthropomorphic figures. Between 1898 and 1910 Clarence Bicknell realized up to 13,000 drawings and reliefs, part of which were then published in [Bic13]. Bicknell identified seven types of figures taking a natural history approach: horned figures (mainly oxen), ploughs, weapons and tools, men, huts and properties, skins and geometrical forms [CB84]. From 1967 Henry de Lumley is in charge of performing research on the site. Figure 2 shows a picture of a bovine engraving, a Bicknell’s relief, and two digitalized reliefs made by de Lumley’s team.



**Figure 2:** A picture of a bovine engraving of the Mont Bego (a), a picture of a bovine relief made by Bicknell on botanic sheet (b), two digitalized reliefs made by de Lumley’s team (c) and (d).

### 2.2. Image Processing of Petroglyphs

The symbol recognition problem is one of the most studied and analyzed research topic in the field of the image processing [LVSM01, TTD06]. But surprisingly, the study of the rock art was touched only minimally by these investigations. Probably, this is due to some unique properties of the petroglyphs (e.g., different petroglyphs may share more

or less the same patterns while being different), which make them unsuitable for recognition tasks.

The work presented in [She80] aimed to catalogue petroglyphs in terms of lengths of parts of animal bodies, and relations among petroglyphs of several regions. In [TTM\*06] Takaki *et al.* proposed new methods to characterize shapes of the petroglyphs and the properties of the group they belong to. In particular, they first extract the skeleton of the petroglyph by applying different image processing algorithms, then their structure is expressed through elementary symbols in order to allow a quantitative comparison. The properties of petroglyph groups are expressed by statistics of simple quantities, such as the numbers of animals and men.

Recently, Zhu *et al.* applied a distance measure based on the Generalized Hough Transform to find meaningful motifs within large collections of rock art images [ZWKL11]. They also proposed a tool called *PetroAnnotator*, which allows human volunteers to “help” computer algorithms segment and annotate petroglyphs [ZWKL09]. Finally, in [SB11] Seidl and Breiteneder proposed a pixel-wise classification for rock art image segmentation and presented some preliminary results.

### 3. An Approach for Petroglyph Classification

One of the most promising approach to achieve low error rates in the classification of images with high variability is the application of flexible matching algorithms [BMP02]. Among them, the deformation models are especially suited to compensate small local changes as they often occur in the presence of image object variability [KDG07]. These models were originally developed for optical character recognition by Keyser *et al.* [KDG07] but it was already observed that it could be applied in other areas such as recognition of medical radiographs [KGN04] and video analysis [DDKN06]. The image distortion model (IDM) yields a distance measure tolerant with respect to local distortions since in the case two images have different values only for a few pixels, due to noise or artifacts irrelevant to classification, the distance between them is compensated by specifying a region in the matching image for each picture element in which it is allowed to detect a best matching pixel.

These properties motivate its use for petroglyph classification. In the following we describe the steps of our petroglyph classification system.

#### 3.1. Shape Normalization

To recognize a petroglyph symbol regardless of its size and position, the input image is normalized to a standard size by translating its center of mass to the origin. The resulting image  $f(x,y)$  is the grid image of the symbol. Then to increase tolerance to local shifts and distortions we smooth and downsample the feature images. In particular, first, to


ensure that small spatial variations in the symbol correspond to gradual changes in the feature values, we apply a Gaussian lowpass filter

$$G(x,y) = e^{-\frac{1}{2}\left(\frac{x^2+y^2}{\sigma^2}\right)}$$

to obtain the smoothed image  $g(x,y)$  according to the following equations

$$g(x,y) = f(x,y) * G(x,y).$$

We then downsample the images by performing symbol removing and resizing (see Figure 3). This further reduces sensitivity to small shifts and improves runtime performance.



**Figure 3:** An example of normalization of the *ax* petroglyph depicted in Fig. 1(d). (a) the image smoothed with the Gaussian filter, (b) the point removed image, and (c) the image resized at  $32 \times 32$  pixels.


#### 3.2. Feature set

To achieve better performances instead of comparing the pixels of the images directly we use derivatives. In particular, for each pixel we consider the horizontal and vertical gradients as features for image matching. Each pixel of a gradient image measures the change in intensity of that same point in the original image, in a given direction. Thus, the horizontal and vertical gradients allow to get the full range of direction.

#### 3.3. Classification

For the classification of petroglyphs we use a deformation model that is robust to distortions and local shifts. In particular, the image deformation model (IDM) performs a pixel-by-pixel value comparison of the query and reference images determining, for each pixel in the query image, the best matching pixel within a region around the corresponding position in the reference image.

The IDM has two parameters: warp range ( $w$ ) and context window size ( $c$ ). Fig. 4 illustrates how the IDM works and the contribution of both parameters, where the warp range  $w$  constrains the set of possible mappings and the  $c \times c$  context window computes the difference between the horizontal and



**Figure 4:** Example of areas affected by the comparison of pixels with IDM, where  $w = 3$  and  $c = 2$ . The query pixel context (indicated by the orange area in the query image) is compared with each equal-sized rectangle within the warping area (dark-green rectangle of the reference image). The warping area is calculated by building a  $m \times m$ , with  $m = (w + c) * 2 + 1$ , square around the corresponding reference pixel (dark-green pixel).

vertical gradient for each mapping. It should be noted that these parameters need to be tuned.

The algorithm requires each pixel in the test image to be mapped to a pixel within the reference image not more than  $w$  pixels from the place it would take in a linear matching. Over all these possible mappings, the best matching pixel is determined using the  $c \times c$  local gradient context window by minimizing the difference with the test image pixel. In particular, the IDM distance  $D$  between two symbols  $S_1$  (the query input) and  $S_2$  (the template) is defined as:

$$D^2 = \sum_{x,y} \min_{d_x, d_y} \|S_1(x + d_x, y + d_y) - S_2(x, y)\|^2$$

where  $d_x$  and  $d_y$  represent pixel shifts and  $S_i(x, y)$  represents the feature values in  $S_i$  from the patch centered at  $x, y$ .

### 3.4. Performance Optimization

One of the limitations of IDM algorithm is the high computational complexity, which is even further increased when the warp range and local context are enlarged. Thus, since applying IDM to all the reference images is too slow, we introduce two optimization strategies to speed up the IDM algorithm.

The first optimization is to prune the set of candidates before applying IDM. We use the simple Euclidean  $L^2$  distance as the pruning metric, and the first  $N$  nearest neighbors found are given in input to the IDM. In particular, we use the distance:

$$D^2 = \sum_{k=1}^K (v_1(k) - v_2(k))^2$$

where  $v_i(k)$  corresponds to the horizontal and vertical gradients of the  $i$ -th image.

The second optimization is the early termination strategy proposed in [SDS08], which relies on the consideration that in  $k$ NN classifiers only the  $k$  nearest neighbors are used in the classification step. Therefore the exact distance of any image with rank greater than  $k$  is not used by the classifier. This means that we can abort the computation of the distance between two reliefs as soon as it exceeds the exact distance of the image with rank  $r$ . Since the latter can only be known after all images in the collection have been compared to the query, we approximate it with the distance of the  $k$  nearest neighbor identified so far.

## 4. Experiments

To demonstrate the validity of the proposed approach an experiment focused on measuring the effectiveness of the IDM algorithm on a real dataset has been performed. The considered dataset was extracted from the image reliefs presented in [dLE09] and it is a representative sample of the Mount Bego petroglyphs. Basically, it contains a number of petroglyph reliefs falling into 10 main classes (anthropomorphic, ax, bull, bullgod, dagger, goddess, oxcart, personage, reticulate, stream). These classes were successively refined, with the help of several archaeologists participating to the project, into 17 classes based on the shape of the petroglyphs and additional information, such as the estimated date of the engravings. The obtained dataset is depicted in Table 3. For each class, we considered three drawings having the specific characteristics of the class. In this way, the overall test set involved into the experiment is composed of 51 drawings.

The 3-fold cross validation test has been successively applied to analyze the performance of the IDM algorithm on the test set. The difficulty to collect data makes particularly suitable this choice. Generally, in  $k$ -fold cross-validation, the original sample is randomly partitioned into  $k$  subsamples. Of the  $k$  subsamples, a single subsample is retained as the validation data for testing the model, and the remaining  $k - 1$  subsamples are used as dataset. The cross-validation process is then repeated  $k$  times (the folds), with each of the  $k$  subsamples used exactly once as the validation data. The  $k$  results from the folds then are averaged to produce a single estimation. In this way all petroglyph reliefs are considered for both dataset and validation, and each relief is used for validation exactly once.

### 4.1. Settings

The best performance of the algorithm was achieved using the following configuration settings:

1. *Image size*, this value was set to 16 pixels. It indicates the size of the reduced image fed to the IDM. We tried also 24 and 32 pixels but no substantial differences have been found. So, we preferred to use the minor size for improving time performance.

2. *Warp range*, together with the local context is one of the parameters of the IDM algorithm. The algorithm obtained the best performance using the value of 3 pixels.
3. *Local context*, this value was set to 2 pixels.

#### 4.2. Results

To evaluate the achieved results we used the  $k$ -nearest-neighbor (kNN) classification method [CH67], which consists in the examination of the first  $k$  results produced by the algorithm.

The overall results of the experiment are listed in Table 1. Data are presented in a 20x9 table where the first row represents the classification method, namely 1-NN for the best matching class, 3-NN and 5-NN for the nearest neighbors methods with 3 and 5 considered hits, respectively. The first column lists the 17 classes considered in the experiment. Basically, each data cell indicates the percentage of times a petroglyph relief is correctly classified by considering the first  $i$  hits of the IDM algorithm result and by applying the classification method (1-NN, 3-NN or 5-NN). As an example, let us consider the  $Ax_B$  row. As for the 1-NN method, during the experiment the  $Ax_B$  images were correctly associated to the class, namely another  $Ax_B$  image appeared as first hit of the IDM algorithm, only in the 33%. In the 67%, it appears either as first or second hit, and finally, it appears in 100% either as first, second, or third hit. If we consider the 3-NN method, the images of the same classes falling into the first three hits are aggregated using the inverse distance weighting. In this way, the class having the highest distance in the new ordered list is suggested to be the class that the query drawing image belongs to. In case of  $Ax_B$ , its images are correctly classified in the 33% considering only the first hit of the result list and in the 100% considering the two best hits. The same analysis has been performed for the 5-NN classification method.

The last row in Table 1 indicates the average values of the columns. Basically, they erase the differences among classes in terms of correct response percentages and report the average behavior of the IDM associated with the different classification methods.

#### 4.3. Discussion

By analyzing the results shown in Table 1 it is possible to notice that the IDM associated to the 1-NN classification method has a precision of 68%, slightly worst in case of 3-NN (61%), and even more worst for 5-NN (54%). Probably, this is due because even though the most similar relief falls into the same class of the query, allowing the 1-NN classifier to correctly recognize, the other hits are not so different from the query but belong to different classes (two symbols for class are in the dataset). In this way, aggregation of the 3-NN and 5-NN makes the choice of the class more difficult

and address it towards wrong classes rather than the class the query belongs to.

Another consideration concerns with the ability of the approach to suggest a number of possible solutions among which to choose the correct class. It is possible to notice that, even if the best classification approach is 1-NN, this is not always true when the algorithm try to suggest a range of possible solutions. Indeed, when considering the most scored two, 3-NN and 5-NN work slightly better than 1-NN (86% and 84% versus 82%). Unfortunately, due to the low number of symbols for class, we cannot extend this consideration for the three or four most scored hits. Anyway, for these cases, the 1-NN allows to correctly classify in the 88% and 92% of cases.

By analyzing the classification rates achieved for each symbol class it is possible to notice some interesting evidences. Among the 51 queries only in four cases the IDM algorithm is not able to retrieve a relief of the same class within the first four hits. As an example, the worst results were achieved for the query Goddess<sub>A</sub>-1 whose classification list was Personage-2, Stream<sub>B</sub>-3, Goddess<sub>D</sub>-2, Goddess<sub>B</sub>-3, and Personage-3. Although the classes are different, overlapping the query to each of these drawings it can be observed that many pixels in the query image match the pixels of the reference images. On the contrary, Goddess<sub>A</sub>-2 and Goddess<sub>A</sub>-3 have a similar structure but the shapes are strongly deformed. This issue could be alleviated by enlarging the warp range and the local context parameters of the algorithm.

Another interesting example may be found by observing the query Dagger<sub>B</sub>-2. In this case the first four hits are: Dagger<sub>A</sub>-3, Dagger<sub>B</sub>-1, Dagger<sub>B</sub>-3, and Dagger<sub>A</sub>-1. Also in this case, the algorithm does not match the correct class in the first hits but allows to identify the classes with images having very similar shapes.

These results highlight the quality of the IDM algorithm in the handling of image deformation but also the natural complexity of the problem faced in this paper. Indeed, in most cases different classes contain similar reliefs and only reasoned-contextual information may help to correctly classify them.

To compare our method to a previously proposed classifier for petroglyph recognition, we apply the Generalized Hough Transform (GHT) algorithm proposed in [ZWKL11] to the considered dataset. This method uses the principle of template matching, which relies on detecting smaller elements matching a template image. In particular, the problem of finding the model's position is transformed to a problem of finding the transformation's parameter that maps the model into the image. GHT uses edge information to define a mapping from orientation of an edge point to a reference point of the shape and computes a measure which rates how well points in the image are likely to be origins of the specified shape.



Symbol	1-NN				3-NN		5-NN	
	1 <sup>st</sup>	2 <sup>st</sup>	3 <sup>st</sup>	4 <sup>st</sup>	1 <sup>st</sup>	2 <sup>st</sup>	1 <sup>st</sup>	2 <sup>st</sup>
Antropomorfe	67	100	100	100	67	100	33	67
Ax <sub>A</sub>	33	67	67	100	33	67	0	100
Ax <sub>B</sub>	33	67	100	100	33	100	33	100
Bull <sub>A</sub>	100	100	100	100	67	100	100	100
Bull <sub>B</sub>	100	100	100	100	100	100	100	100
Bullgod	67	100	100	100	67	100	67	100
Dagger <sub>A</sub>	100	100	100	100	100	100	100	100
Dagger <sub>B</sub>	33	67	67	100	67	67	33	67
Goddess <sub>A</sub>	67	67	67	67	33	67	33	67
Goddess <sub>B</sub>	67	67	67	67	67	67	67	67
Goddess <sub>C</sub>	67	67	100	100	67	67	67	67
Goddess <sub>D</sub>	67	67	67	67	33	67	0	67
Oxcart	100	100	100	100	67	100	67	100
Personage	33	67	67	67	33	67	0	67
Reticulate	67	100	100	100	67	100	67	100
Stream <sub>A</sub>	100	100	100	100	100	100	100	100
Stream <sub>B</sub>	67	67	100	100	33	100	67	67
<b>total</b>	<b>68</b>	<b>82</b>	<b>88</b>	<b>92</b>	<b>61</b>	<b>86</b>	<b>54</b>	<b>84</b>

**Table 1:** Classification rates of the proposed IDM algorithm.

	1-NN		3-NN		5-NN	
	IDM	GHT	IDM	GHT	IDM	GHT
1 <sup>st</sup>	68%	35%	61%	31%	54%	27%
2 <sup>st</sup>	82%	43%	86%	47%	84%	49%
3 <sup>st</sup>	88%	47%	-	-	-	-
4 <sup>st</sup>	92%	53%	-	-	-	-

**Table 2:** Comparison of the classification rates achieved with IDM (images normalized at 16 pixels) and GHT (images normalized at 32 pixels).

The same evaluation procedure as for the IDM algorithm has been applied. The results shown in Table 2 demonstrate that our approach outperforms the method in [ZWKL11]. In particular, the GHT algorithm with a symbol normalization at 32 pixels achieves, in a time comparable to IDM algorithm, classification rates which are almost halved with respect to IDM. The performance of the GHT algorithm is highly dependent on the results from the edge detector, so the input image must be carefully chosen for greater edge detection. In fact, the application of GHT on noisy images reduces the efficiency of the algorithm due to some edge points being missed because they were not defined as edges. Therefore, the proposed IDM algorithm provides a more robust similarity measure for petroglyph reliefs.

## 5. Conclusions and Future Work

In this paper we presented a classifier for petroglyph symbol reliefs robust to distortions and local shifts. The method is based on IDM algorithm [KDGNO7] to measure the distances between the queries and the reference images, which is an effective means of compensating for small local im-

age. The experimental results show the potential of the proposed method for petroglyph classification with a classification rate of 68%, which considerably improves a previous distance proposed for petroglyph recognition of about 33% [ZWKL11]. These results are achieved on a representative dataset of Mount Bego petroglyphs, which includes all the main petroglyph classification challenges.

With respect to previous works where the application of the IDM algorithm showed very good performances (generally the correct rate was higher than 97%) [KGN04, DWK\*05], with petroglyph symbols the IDM performance dramatically decreases (68% for 1-NN, 61% for 3-NN, 54% for 5-NN) highlighting the main difficulties to manage this kind of datasets. Results improve when the approach is used for suggesting a set of possible classes. In this cases, when we classify using 1-NN, the rate of correct responses is higher than 80% by considering the first two petroglyph classes, close to 90% for the first three, and higher 90% for the first four. By using 3-NN and 5-NN the classification slightly enhances considering the first two classes, i.e., 86% and 84%. By analyzing the wrong classifications it is possible to notice that the errors are mainly due to the high intra-class variability and low inter-class variability of the dataset, and well-reasoned contextual information may help to correctly classify the ambiguous engravings. These challenges make the 3-NN and 5-NN classifier more error-prone. Indeed, in some cases, even though the best scored image belongs to the same class of the query, it might happen that similar images of other classes obtain better scores of the images of the query class. When this occurs, the aggregation of the results can change the order of the classes introducing errors.

The main research challenges for the future will be the investigation of other optimization strategies for IDM and the validation of the results on a larger dataset. In particular, in order to improve the efficiency of the approach we are going to explore the possibility to reduce the number of IDM comparisons by means of a clustering algorithm. Moreover, we intend to investigate the use of a distance with a behavior orthogonal to IDM, i.e., able to achieve good performances for IDM misclassified images. In this way, by applying a suitable weighting scheme it would be possible to exploit the advantages of both techniques.

We are also interested in the application of petroglyph recognition algorithms for integrating the image reliefs made by Bicknell with those of de Lumley. Indeed, even if the Bicknell legacy is relatively small, a manual integration would be time-consuming and error-prone [PQM\*11]. Therefore, the proposed algorithm can help archaeologists integrate the data in a semi-automatic way.

Finally, we intend to investigate the use of query by sketch as a technique to ease user interaction and improve retrieval effectiveness in the repository [CDM\*08].

## 6. Acknowledgments

This research is supported by the “Indiana MAS and the Digital Preservation of Rock Carvings: A multi-agent system for drawing and natural language understanding aimed at preserving rock carving” FIRB project funded by the Italian Ministry for Education, University and Research, under grant RBFR10PEIT.

## References

- [Bic13] BICKNELL C. M.: *A Guide to the Prehistoric Rock Engravings in the Italian Maritime Alps*. printed by Giuseppe Bessone, 1913. 2
- [BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 4 (Apr. 2002), 509–522. 3
- [CB84] CHIPPINDALE C., BICKNELL C.: Archaeology and science in the 19th century. *Antiquity* 58, 224 (Sept. 1984), 185–193. 2
- [CDM\*08] CASELLA G., DEUFEMIA V., MASCARDI V., COSTAGLIOLA G., MARTELLI M.: An agent-based framework for sketched symbol interpretation. *J. Vis. Lang. Comput.* 19, 2 (2008), 225–257. 7
- [CH67] COVER T., HART P.: Nearest neighbor pattern classification. *IEEE Transactions on Information Theory* 13, 1 (1967), 21–27. 5
- [DDKN06] DREUW P., DESELAERS T., KEYSERS D., NEY H.: Modeling image variability in appearance-based gesture recognition. In *ECCV Workshop on Statistical Methods in Multi-Image and Video Processing* (Graz, Austria, May 2006), pp. 7–18. 3
- [dLE09] DE LUMLEY H., ECHASSOUX A.: The rock carvings of the chalcolithic and ancient bronze age from the mont bego area. The cosmogonic myths of the early metallurgic settlers in the southern alps. *L'Anthropologie* 113, 5P2 (2009), 969–1004. 1, 2, 4
- [DPT\*12] DEUFEMIA V., PAOLINO L., TORTORA G., TRAVERSO A., MASCARDI V., ANCONA M., MARTELLI M., BIANCHI N., DE LUMLEY H.: Investigative analysis across documents and drawings: visual analytics for archaeologists. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (2012), ACM, pp. 539–546. 1
- [DWK\*05] DESELAERS T., WEYAND T., KEYSERS D., MACHEREY W., NEY H.: FIRE in ImageCLEF 2005: Combining content-based image retrieval with textual information retrieval. In *CLEF* (2005), Peters C., Gey F. C., Gonzalo J., Müller H., Jones G. J. F., Kluck M., Magnini B., de Rijke M., (Eds.), vol. 4022 of *Lecture Notes in Computer Science*, Springer, pp. 652–661. 1, 6
- [KDG07] KEYSERS D., DESELAERS T., GOLLAN C., NEY H.: Deformation models for image recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29 (2007), 1422–1435. 1, 3, 6
- [KGN04] KEYSERS D., GOLLAN C., NEY H.: Local context in non-linear deformation models for handwritten character recognition. In *ICPR* (4) (2004), pp. 511–514. 1, 3, 6
- [LVSM01] LLADÓS J., VALVENY E., SÁNCHEZ G., MARTÍ E.: Symbol recognition: Current advances and perspectives. In *GREC* (2001), Blostein D., Kwon Y.-B., (Eds.), vol. 2390 of *Lecture Notes in Computer Science*, Springer, pp. 104–127. 2
- [MDM\*12] MASCARDI V., DEUFEMIA V., MALAFRONTÉ D., RICCIARELLI A., BIANCHI N., DE LUMLEY H.: Rock art interpretation within indiania mas. In *Agent and Multi-Agent Systems. Technologies and Applications*, Jezic G., Kusek M., Nguyen N.-T., Howlett R., Jain L., (Eds.), vol. 7327 of *Lecture Notes in Computer Science*. Springer, 2012, pp. 271–281. 1
- [PQM\*11] PAPALEO L., QUERCINI G., MASCARDI V., ANCONA M., TRAVERSO A., DE LUMLEY H.: Agents and ontologies for understanding and preserving the rock art of Mount Bego. In *ICAART* (2) (2011), Filipe J., Fred A. L. N., (Eds.), SciTePress, pp. 288–295. 7
- [SB11] SEIDL M., BREITENEDER C.: Detection and Classification of Petroglyphs in Gigapixel Images – Preliminary Results. In *Proceedings of the 12th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (2011), Eurographics Association, pp. 45–48. 3
- [SDS08] SPRINGMANN M., DANDER A., SCHULTD H.: Improving efficiency and effectiveness of the image distortion model. *Pattern Recognition Letters* 29, 15 (2008), 2018–2024. 4
- [She80] SHER Y. A.: *Petroglyphs in Central Asia*. Nauka, 1980. 2
- [TTD06] TOMBRE K., TABBONE S., DOSCH P.: Musings on symbol recognition. In *Graphics Recognition. Ten Years Review and Future Perspectives*, Liu W., Lladós J., (Eds.), vol. 3926 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2006, pp. 23–34. 2
- [TTM\*06] TAKAKI R., TORIWAKI J., MIZUNO S., IZUHARA R., KHUDJANAZAROV M., REUTOVA M.: Shape analysis of petroglyphs in central asia. *Forma* 21 (2006), 91–127. 3
- [ZWKL09] ZHU Q., WANG X., KEOGH E., LEE S.-H.: Augmenting the generalized hough transform to enable the mining of petroglyphs. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2009), ACM, pp. 1057–1066. 3
- [ZWKL11] ZHU Q., WANG X., KEOGH E., LEE S.-H.: An efficient and effective similarity measure to enable data mining of petroglyphs. *Data Mining and Knowledge Discovery* 23 (2011), 91–127. 1, 2, 3, 5, 6




















































Class	Relief Images			Class	Relief Images		
Antropomorphe				Goddess <sub>B</sub>			
Ax <sub>A</sub>				Goddess <sub>C</sub>			
Ax <sub>B</sub>				Goddess <sub>D</sub>			
Bull <sub>A</sub>				Oxcart			
Bull <sub>B</sub>				Personage			
Bullgod				Reticulate			
Dagger <sub>A</sub>				Stream <sub>A</sub>			
Dagger <sub>B</sub>				Stream <sub>B</sub>			
Goddess <sub>A</sub>							

Table 3: The dataset of petroglyph reliefs used for evaluation.