Analyzing and Visualizing Ancient Maya Hieroglyphics Using Shape: from Computer Vision to Digital Humanities

Rui Hu
Idiap Research Institute, Switzerland

Carlos Pallán Gayol
University of Bonn, Germany

Jean-Marc Odobez
Idiap Research Institute
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Daniel Gatica-Perez
Idiap Research Institute
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Abstract

Maya hieroglyphic analysis requires epigraphers to spend a significant amount of time browsing existing catalogs to identify individual glyphs. Automatic Maya glyph analysis provides an efficient way to assist scholars' daily work. We introduce the Histogram of Orientation Shape Context (HOOSC) shape descriptor to the digital humanities community. We discuss key issues for practitioners and study the effect that certain parameters have on the performance of the descriptor. Different HOOSC parameters are tested in an automatic ancient Maya hieroglyph retrieval system with two different settings, namely when shape alone is considered and when glyph co-occurrence information is incorporated. Additionally, we developed a graph-based glyph visualization interface to facilitate efficient exploration and analysis of hieroglyphs. Specifically, a force-directed graph prototype is applied to visualize Maya glyphs based on their visual similarity. Each node in the graph represents a glyph image; the width of an edge indicates the visual similarity between the two according glyphs. The HOOSC descriptor is used to represent glyph shape, based on which pairwise glyph similarity scores are computed. In order to evaluate our tool, we designed evaluation tasks and questionnaires for two separate user groups, namely a general public user group and an epigrapher scholar group. Evaluation results and feedback from both groups show that our tool provides intuitive access to explore and discover the Maya hieroglyphic writing, and could potentially facilitate epigraphy work. The

positive evaluation results and feedback further hints the practical value of the HOOSC descriptor.

1. Introduction

Technological advances in digitization, automatic image analysis, and information management are enabling the possibility to analyze, organize and visualize large cultural datasets. As one of the key visual cues, shape has been used in various image analysis tasks such as handwritten character recognition (Fischer et al., 2012; Franken et al., 2013), and sketch analysis (Eitz et al., 2012). We assess a shape descriptor, within the application domain of Maya hieroglyphic analysis. Our aim is to introduce this descriptor to the wider Digital Humanities (DH) community, as a shape analysis tool for DH-related applications. Two example application systems, namely an automatic glyph retrieval framework and an interactive glyph visualization interface are presented.

The Maya civilization is one of the major cultural developments in ancient Mesoamerica. The ancient Maya language created uniquely pictorial forms of hieroglyphic writing (Stone et al., 2011). Most Maya texts were written during the Classic period (AD 250-900) of the Maya civilization on various media types, including stone monuments, architectural elements, as well as personal items. A special class of Maya texts was written on bark cloths, made from the inner bark of certain trees (the main being the ficus tree), as folding books from the Post-Classic period (AD 1000-1519). Only three such books (namely the Dresden, Madrid and Paris codices) are known to have survived the Spanish Conquest. A typical Maya codex page contains icons, main sign glyph blocks, captions, calendric signs, etc. In this paper, we are interested in the main signs.

Maya hieroglyphic analysis requires epigraphers to spend a significant amount of time browsing existing catalogs to identify individual glyphs. Automatic Maya glyph analysis provides a potentially efficient way to assist scholars' daily work. Identifying unknown hieroglyphs can be addressed as a shape matching problem. A robust shape descriptor called Histogram of Orientation Shape Context (HOOSC) was developed in (Roman-Rangel et al., 2011). Since then, HOOSC has been successfully applied for automatic analysis of other cultural heritage data, such as Oracle-Bones Inscriptions of ancient Chinese characters (Roman-Rangel, 2012), and ancient Egyptian hieroglyphs (Franken et al., 2013). It has also been applied for generic sketch and shape image retrieval (Roman-Rangel, 2012). Our recent work extracted a statistic Maya language model and incorporated it for glyph retrieval (Hu et al., 2015).

In another direction, data visualization techniques, which organize and present data in a structured graphical format, can be applied to visualize glyph images, enable efficient browsing and search of ancient hieroglyph datasets. Such systems can enable mining and discovery of visual and semantic patterns of the ancient Maya writing, which could potentially facilitate research aimed at advancing hieroglyphic decipherment. We developed an interactive, graph-based glyph visualization interface.

In this paper, we introduce the HOOSC descriptor to the digital humanities community. We discuss key issues for practitioners and study the effect that certain parameters have on the performance of the descriptor. Different HOOSC parameters are tested in an automatic ancient Maya hieroglyph retrieval system with two different settings, namely when shape alone is considered and when glyph co-occurrence information is incorporated. Additionally, we present a graph-based visualization interface, where each node in the graph represents a glyph image, an edge between two nodes indicate the visual similarity of the two corresponding glyphs. Glyph visual similarity measurement is an important factor which affects the layout of the graph. In this paper, the HOOSC descriptor was used to represent the shape of each glyph, based on which glyph visual similarity scores were computed. Our interface was developed using the force directed graph prototype of the D3 (Data-Driven Document) web-based visualization approach (Bostock et al., 2011).

The goal of this paper is three-fold:

1. Introduce the HOOSC descriptor to be used in DH-related shape analysis tasks (code available at: http://www.idiap.ch/project/maaya/code/);
2. Discuss key issues for practitioners, namely the effect that certain parameters have on the performance of the descriptor. We describe the impact of such choices on different data types, specially for 'noisy' data as it is often the case with DH image sources;
3. Present and evaluate a graph-based, glyph visualization interface to explore Maya hieroglyph data. The HOOSC descriptor is applied to represent the shape feature of glyph images, based on which pairwise glyph visual similarity scores are computed to define the edges in the graph.

A preliminary version of this paper was presented in the Digital Humanities conference (Hu et al., 2016).

The rest of the paper is organized as follows: In section 2, we introduce the HOOSC descriptor and discuss the key issues of certain parameters have on the performance of the descriptor in a

state-of-the-art glyph retrieval system; in section 3, we present a graph-based interactive glyph visualization interface, where the HOOSC descriptor is used to represent the shape of glyph images; we draw concluding remarks in section 4.

2. Automatic Maya Hieroglyph Recognition

We conduct glyph recognition with a retrieval system proposed in (Hu et al., 2015). Unknown glyphs are considered as queries to match with a database of known glyphs (retrieval database). Shape and context information are considered. Fig. 1 illustrates a schema of our approach. We study the effect of different HOOSC parameter choices on the retrieval results.

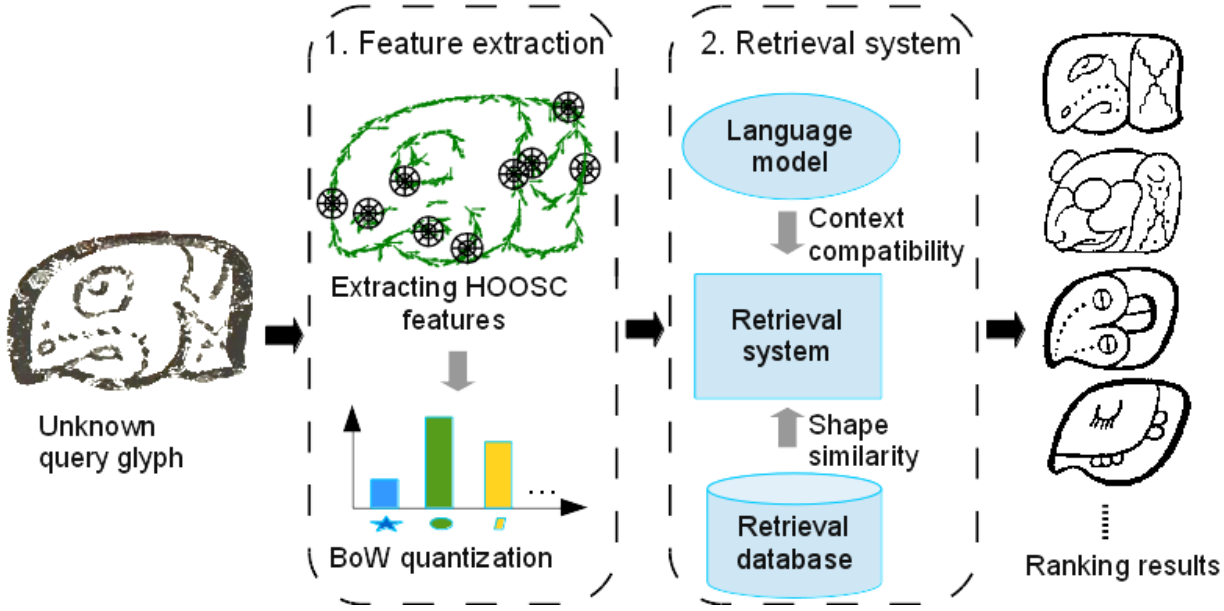

Fig. 1 Retrieval system pipeline.

2.1 Datasets

We use three datasets, namely the 'Codex', 'Monument' and 'Thompson'. The first two are used as queries to search within the retrieval database ('Thompson').

Fig. 2 Digitization quality: (left) raw glyph blocks cropped from Dresden codex; (middle) clean raster images produced by removing the background noise; (right) reconstructed high-quality vectorial images.

The 'Codex' dataset contains glyph blocks from the three surviving Maya codices. Fig. 2 shows examples of three raw glyph block images cropped from the ancient codices (Fig. 2 left); their clean raster versions (Fig. 2 middle); and high-quality reconstructed vectorial images (Fig. 2 right). Clean raster and reconstructed versions are manually generated by epigraphers in our team. The clean raster images are produced by manually removing the background area from the raw images; whereas the reconstructed forms are generated by further carefully reconstructing the broken lines and missing strokes. Glyph blocks are typically composed of combinations of individual signs. Fig. 3 shows individual glyphs segmented from blocks in Fig. 2. Note the different degradation levels across samples. We use two sub-datasets: 'codex-small', composed of 156 glyphs segmented from 66 blocks, for which we have both clean raster and high-quality reconstructed vectorial representations (see Fig. 3) to study the impact of the different data qualities on the descriptor; and a 'codex-large' dataset, which is more extensive, comprising only the raster representation of 600 glyphs from 229 blocks.

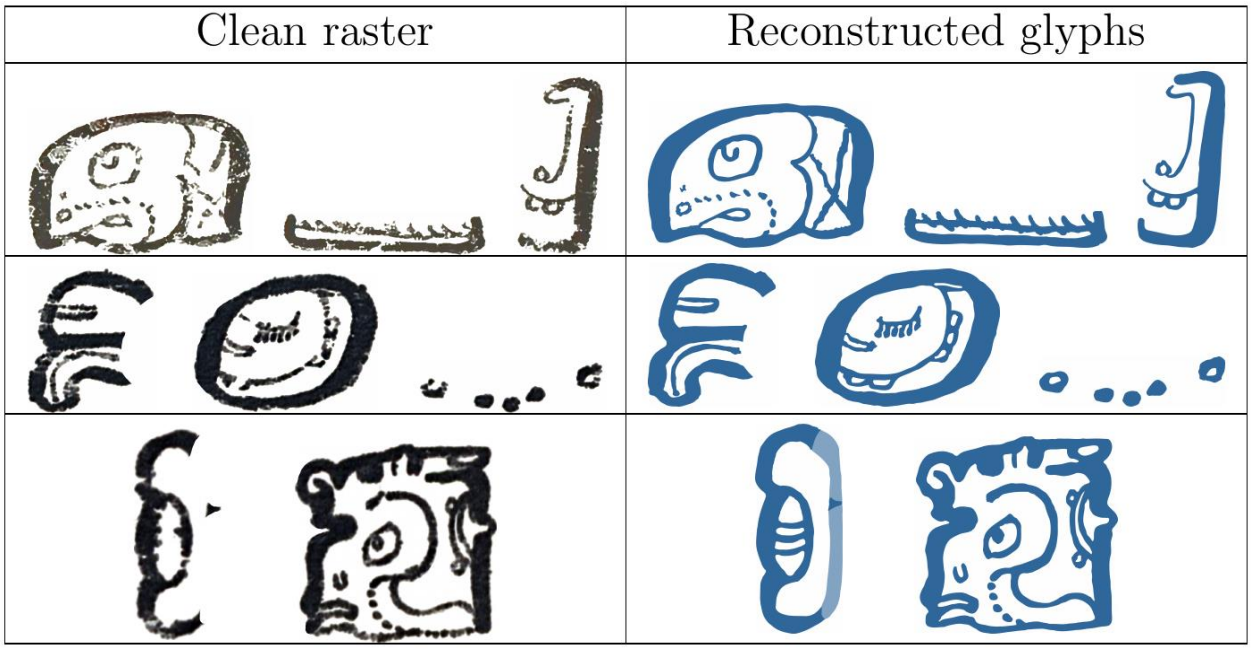

Fig. 3 Example glyph strings generated from blocks shown in Figure 2.

The 'Monument' dataset is an adapted version of the Syllabic Maya dataset used in (Roman-Rangel et al., 2011), which contains 127 glyphs of 40 blocks extracted from stone monuments. It is a quite different data source to the codex data, in terms of historical period, media type, and data generation process. Samples are shown in Fig. 4.

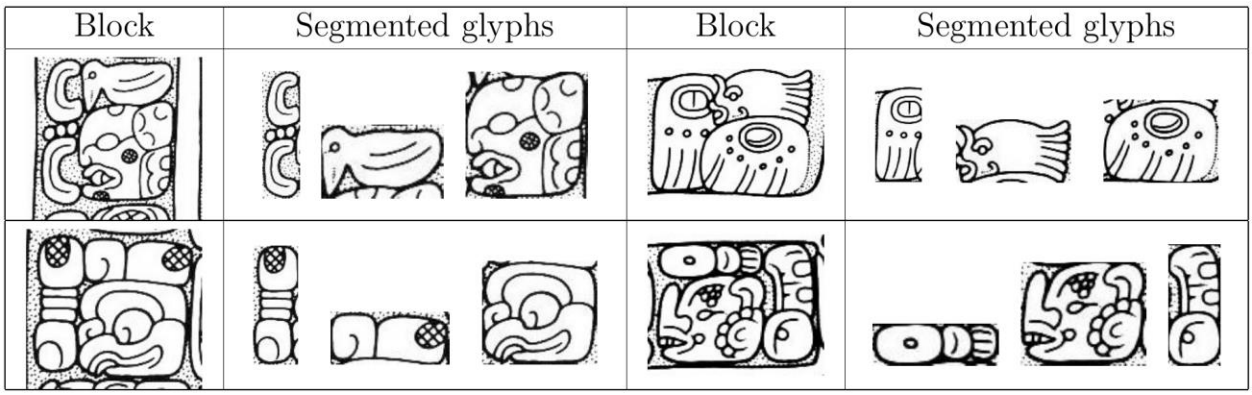

Fig. 4 Example blocks and segmented glyph strings form the 'Monument' dataset.

To form the retrieval database ('Thompson'), we scanned and segmented all the glyphs from the Thompson catalog (Thompson, 1962). The database contains 1487 glyph examples of 892 different sign categories. Each category is usually represented by a single example image. Sometimes multiple examples are included; each illustrates a different visual instance or a rotation variant. Fig. 5 shows glyph examples.

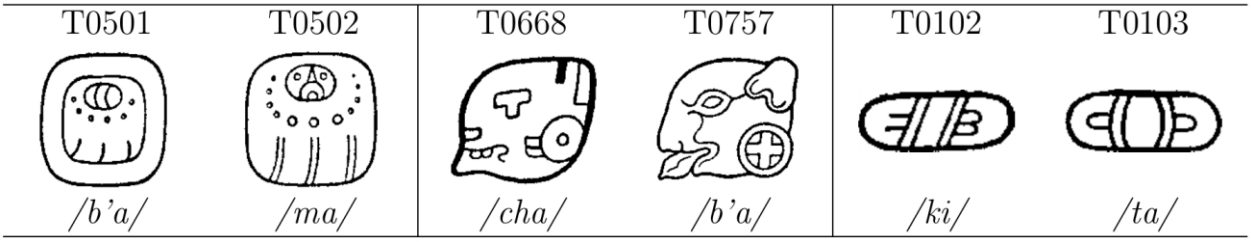

Fig. 5 Thompson numbers, visual examples, and the syllabic values of glyph pairs. Each pair contains two different signs with similar visual features (Hu et al., 2015). All examples are taken from (Thompson, 1962).

2.2 Shape-based retrieval

Feature extraction and similarity matching are the two main steps for our shape-based glyph retrieval framework.

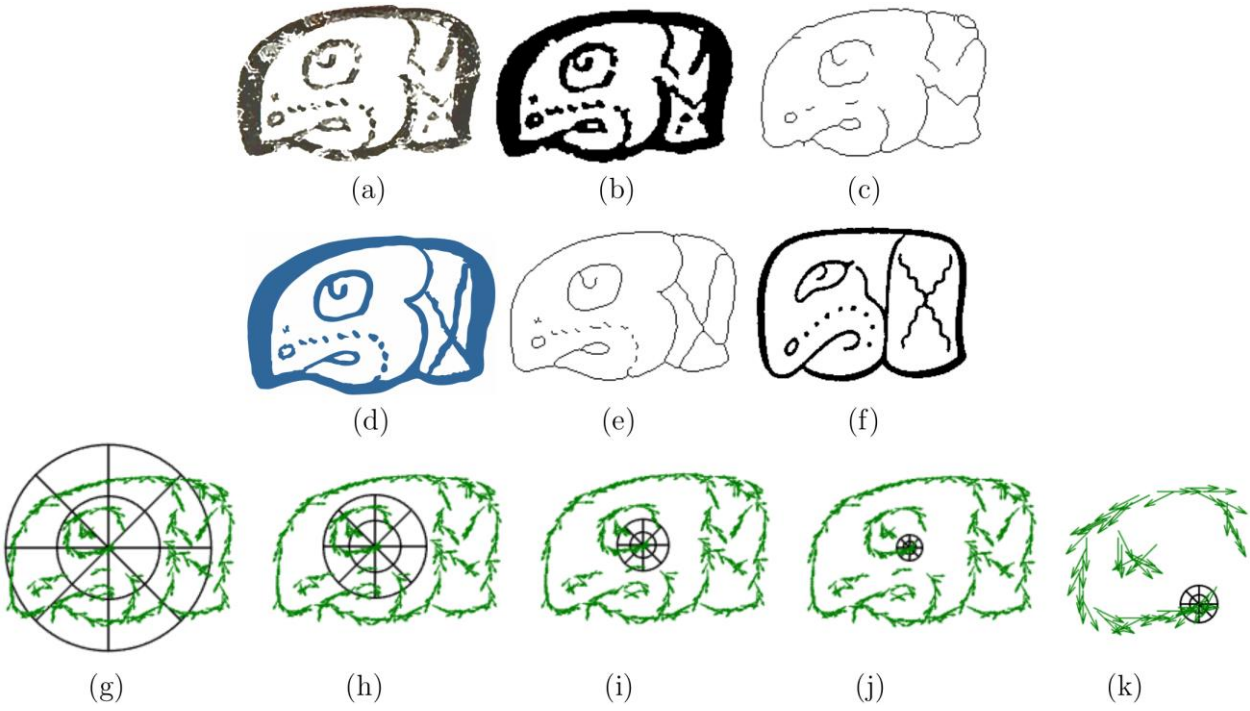

Fig. 6 Extracting HOOSC descriptor: (a) input clean raster image; (b) binary image; (c) thinned edge of (b); (d) reconstructed vector representation of (a); (e) thinned edge of (d); (f) corresponding groundtruth image in the catalog; (g)-(k) spatial partition of a same pivot point with five different ring sizes (1, ½, ¼, 1/8, 1/16, all defined as a proportion to the mean of the pairwise distance between pivot points) on the local orientation field of the thinned edge image (c). Note that we zoomed in to show the spatial context of 1/16 in (k).

Glyphs are first pre-processed into thin lines. To do so, an input glyph (Fig. 6(a)) is first converted into a binary shape (Fig. 6 (b)). Thin lines (Fig. 6(c)) are then extracted through

mathematical morphology operations. Fig. 6(d)-(e) show the high-quality reconstructed binary image and the extracted thin lines.

HOOSC descriptors are then computed at a subset of uniformly sampled pivot points along the thin lines. HOOSC combines the strength of Histogram of Orientation Gradient (Dalal et al., 2005) with circular split binning from the shape context descriptor (Belongie et al., 2002). Given a pivot point, the HOOSC is computed on a local circular space centred at the pivot's location, partitioned into rings and evenly distributed angles. Fig. 6 (g)-(k) show different sizes of the circular space (referred to as spatial context) partitioned into 2 rings and 8 orientations. A Histogram-of-orientation-gradient is calculated within each region. The HOOSC descriptor for a given pivot is the concatenation of histograms of all partitioned regions.

We then follow the Bag-of-Words (BoW) approach, where descriptors are quantized as visual words based on the vocabulary obtained through K-means clustering on the set of descriptors extracted from the retrieval database. A histogram representing the count of each visual word is then computed as a global descriptor for each glyph. In all experiments, we use vocabulary size k=5000.

Each query is matched with glyphs in the retrieval database, by computing shape feature similarity using the L1 norm distance.

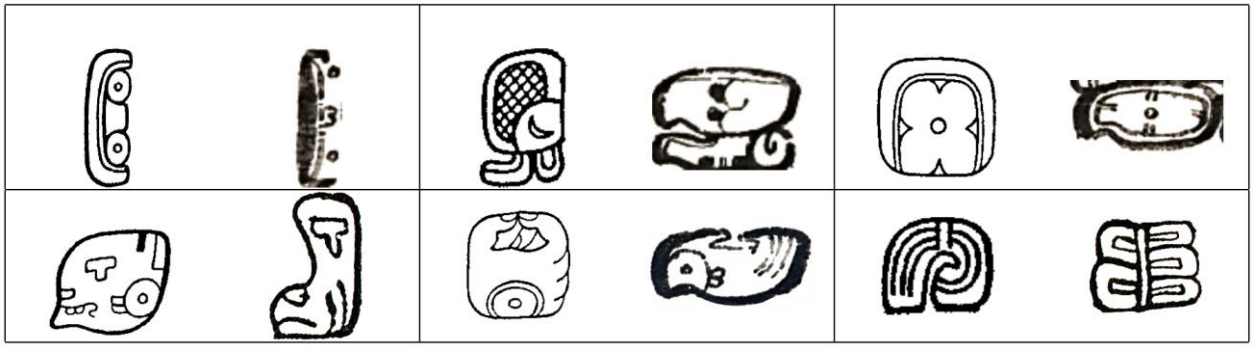

Fig. 7 Six pairs of glyph signs (Hu et al., 2015). Each pair contains a query glyph from the 'Codex' dataset (right) and their corresponding groundtruth in the catalog (left).

2.3 Incorporating context information

Shape alone is often ambiguous to represent and distinguish between images. In the case of our data, different signs often share similar visual features (see Fig. 5); glyphs of the same sign category vary with time, location, and the different artists who produced them (see Fig. 7);

additionally, surviving historical scripts often lose visual quality over time. Context information can be used to complement the visual features.

Glyph co-occurrence within single blocks encodes valuable context information. To utilize this information, we arrange glyphs within a single block into a linear string according to the reading order (see Fig. 3 and Fig. 4) and consider the co-occurrence of neighbouring glyphs using an analogy to a statistical language model. For each unknown glyph in the string, we compute its probability to be labelled as a given category by considering not only the shape similarity, but also the compatibility to the rest of the string.

We apply the two glyph co-occurrence models (statistic language models) extracted in (Hu et al., 2015), namely the ones derived from the Maya Codices Database (Vail et al., 2013) and the Thompson catalog (Thompson, 1962), which we refer to as the 'Vail' and the 'Thompson' models. We use Vail model with smoothing factor $\alpha=0$ for the 'Codex' data, and the Thompson model with $\alpha=0.2$ for the 'Monument' data, which have shown to perform well in (Hu et al., 2015).

2.4 Experiments and Results

Our aim is to demonstrate the effect of various HOOSC parameters on retrieval results.

2.4.1 Experimental setting

We illustrate the effect of 3 key parameters:

- Size of the spatial context region within which HOOSC is computed

A larger region encodes more context information and therefore captures more global structure of the shape. However, in the case of image degradation, a larger region could contain more noise. We evaluate five different spatial contexts as shown in Fig. 6(g)-(k). The circular space is distributed over 8 angular intervals.

- Number of rings to partition the local circular region

This parameter represents different partition details. We evaluate either 1 or 2 rings, the inner ring covers half the distance to the outer ring. Each region is further characterized by a 8-bin histogram of the local orientations.

- Position information

Relative position (i, j) of a pivot in the 2-D image plane can be concatenated to the corresponding HOOSC feature.

2.4.2 Results and discussion

Fig. 8 shows the average groundtruth ranking in the retrieval results with different parameter settings, on three query sets, *e.g.* 'Codex-large', 'Codex-small' and 'Monument'. Each query image usually has only one correct match (groundtruth) in the retrieval database. The smaller the average ranking value, the better the result. From Fig. 8 we can see the following:

- In most cases, the best results are achieved by using the largest spatial context, with finer partitioning details (2 rings in our case);
- When the location information is not considered, results show a general trend of improving with increasing ring sizes. However, the results are more stable when the position information is encoded, *e.g.* a smaller ring size can also achieve promising results when the location information is incorporated. This is particularly useful when dealing with noisy data, where a smaller ring size is preferred to avoid extra noise been introduced by a larger spatial context;
- The results do not benefit from a finer partition when a small spatial context is considered. However, results improve with finer partitions when the spatial context becomes larger.
- Position information is more helpful when a small spatial context is considered.

Fig. 9 shows example query glyphs and their top returned results.

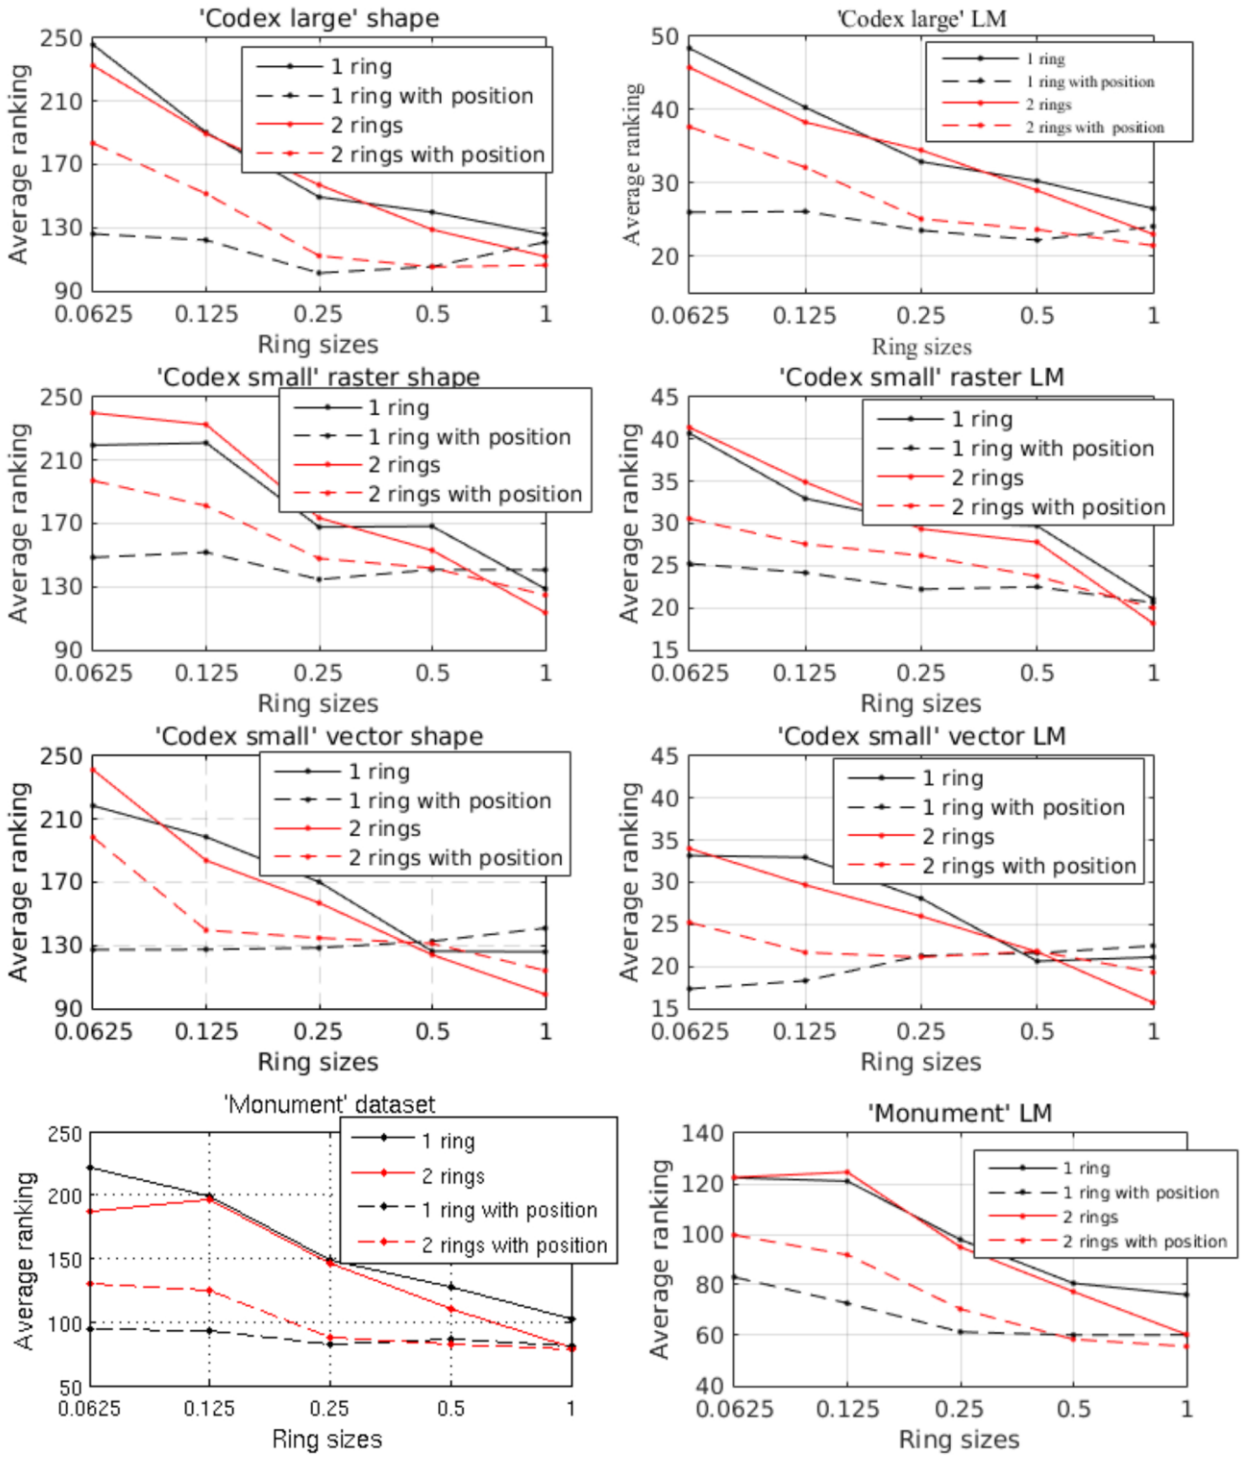

225

226 *Fig. 8 Retrieval results on each dataset, with various feature representation choices. (left) shape-*
227 *base results; (right) incorporating glyph co-occurrence information.*

Fig. 9 Example queries (first column) and their top returned retrieval results, ranked from left to right in each row. Groundtruth images are highlighted in green bounding boxes.

3. Graph-based Glyph Visualization Interface

We developed a web-based visualization interface to enable efficient exploration of ancient Maya hieroglyph signs. Our tool can not only be used by general public users to browse and explore the ancient Maya hieroglyph writing system, but also by epigraphy scholars as a tool to facilitate their daily work of glyph analysis, as well as to discover visual and semantic glyph patterns.

In this section, we first explain the methodology that we followed to develop this tool and its functionalities; we then present our evaluation tasks, together with the evaluation results and feedback collected from both general users and epigraphy expert users.

3.1 Visualization approach

Maya hieroglyph sign examples, scanned and segmented from the Thompson catalog (Thompson, 1962), are illustrated in a force-directed graph. Each node in the graph corresponds to one glyph image in the database. The width of an edge indicates the visual similarity score between the two corresponding glyphs. A detailed explanation of how to compute the glyph similarity scores, and the specific design and functionalities of our tool are introduced below.

3.1.1 Glyph visual similarity measurement

Visual similarity is an important factor which defines the layout of the graph, and therefore affects the performance of the visualization tool. We use the HOOSC descriptor to represent glyph shape, based on which the visual similarity scores are computed.

Given the glyph sign examples in the Thompson catalog, we focus on the signs that appear in the three surviving ancient Maya codices. Based on our annotation of the Maya codex database, there are in total 288 different glyph sign categories appeared in the three codices, which corresponds to 657 image examples in the Thompson catalog. Each of these images defines a node in the graph.

Following the framework introduced in section 2.2, we first extract the HOOSC descriptors for each image and then compute glyph similarity scores between image pairs. Specifically, we conduct the following process. First, each glyph image is pre-processed into thin lines; 400 evenly distributed points along the lines are randomly selected as pivot points. The HOOSC descriptor for each pivot point is then computed on a circular region centered at the pivot location, with a diameter equal to the mean of the pairwise distance between pivot points, which is partitioned into 2 rings and 8 equally distributed orientations. The bag of words pipeline, with

vocabulary size k=5000, is then applied to compute a global histogram representation indicating the count of each visual word. Finally, the pairwise glyph visual similarity is computed based on the L1 norm distance measure.

The resulted graph has 657 nodes and 215496 edges. In order to reduce the complexity of the graph, we reduced the number of edges by applying a threshold to the similarity scores. Two glyph images are only considered to be similar if their visual similarity score is larger than a pre-defined threshold α, and therefore their corresponding nodes in the graph are connected. Otherwise, the two glyphs are not considered to be sufficiently similar, and their corresponding nodes in the graph are not connected. A higher threshold indicates that fewer glyph pairs are considered to be visually similar, which results in fewer edges in the graph. In contrast, a smaller threshold indicates that more glyph pairs are considered to be visually similar, which leads to a more connected graph (more edges). In this paper, we take the threshold α=0.2 as an example for user evaluation tasks in the following sections. The reduced graph has 687 edges. Visualization graph with different threshold scores can be explored at: http://lab.idiap.ch/maaya/.

3.1.2 Design and functionality

We use the force-directed graph layout provided in the D3 visualization approach. D3 is a web-based, representation-transparent visualization approach. Given a glyph image database, we initialize the graph with random node positions and the pre-computed pairwise glyph similarity scores as edge weights. The force-directed graph optimizes to position nodes of the graph so that there are as few crossing edges as possible, and all the edges are of roughly similar length. As a result, nodes in the graph often fall in several groups, each containing glyphs of visually similar patterns. Fig. 10 illustrates our visualization interface. It can be seen that glyphs are grouped into visually similar clusters, such as vertical, horizontal, squared, and knot shaped patterns. This could naturally function as a shape-based glyph indexing system, which could help users quickly identify an unknown glyph based on its visual features.

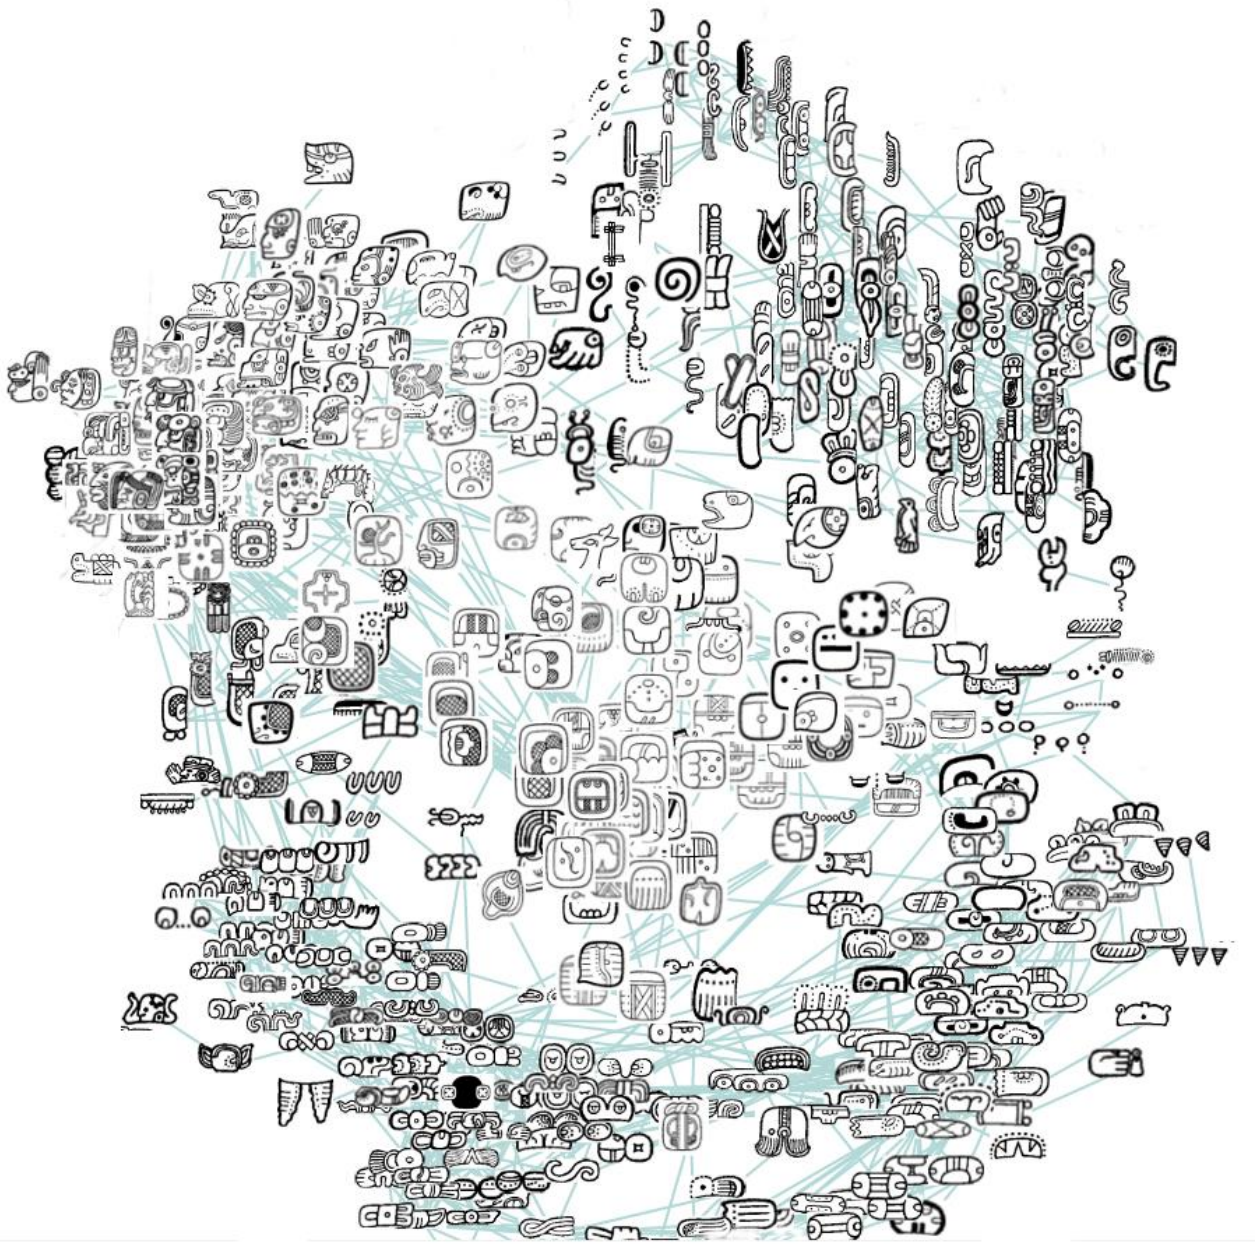

Fig. 10 Graph-based visualization interface.

We adapt a similar design to the 'movie network' example in D3[1]. Our visualization tool allows zooming, dragging, and clicking actions to facilitate exploring and analysis of the glyph database. Upon clicking on a node, an information panel pops up. It shows the example image of the corresponding node in the graph and more detailed information of the individual glyph. Fig. 11 shows the pop up information panel of an example glyph. Three types of information are

[1] Movie network example: http://bl.ocks.org/paulovn/9686202

included in the panel. The first type shows the basic glyph information, which includes the glyph example image shown at the top of the panel, the Thompson code ('Glyph category'), and the index of this glyph variant ('Image name'). The second type provides information about external resources, including the glyphs' 'Appearance in codices' and 'Phonetic and semantic values'. Finally, the third type presents information that is automatically extracted from images based on shape analysis ('Similar to') and the statistic glyph co-occurrence information ('Co-occur with') extracted from the three surviving ancient Maya codices. Below, we give detailed explanation of the latter two types of information shown in the pop up window.

Links to external resources.

By clicking on glyphs in the list of 'Appearance in codices', our interface leads users to an online database provided by the Maya codices database project (Vail et al., 2013), where users can browse through all the codex 'frames' from the three ancient codices in which this glyph appears.

By clicking on glyphs in the list of 'Phonetic and semantic values', the interface conducts users to the John Montgomery Dictionary of Maya Hieroglyphs (Montgomery et al., 2007) via the Foundation for the Advancement of Mesoamerican Studies, Inc. (FAMSI) webpage (http://www.famsi.org/), which gives more detailed information of the given glyph, including its pronunciation and semantic meaning if applicable, decoded by epigraphers over the years.

These functions provide users with access to valuable and already existing Maya hieroglyphic resources, which can be used to facilitate scholars' work, and also to help members of the general public interested in exploring the ancient Maya scripts. Our interface here serves as a connecting point to richer resources.

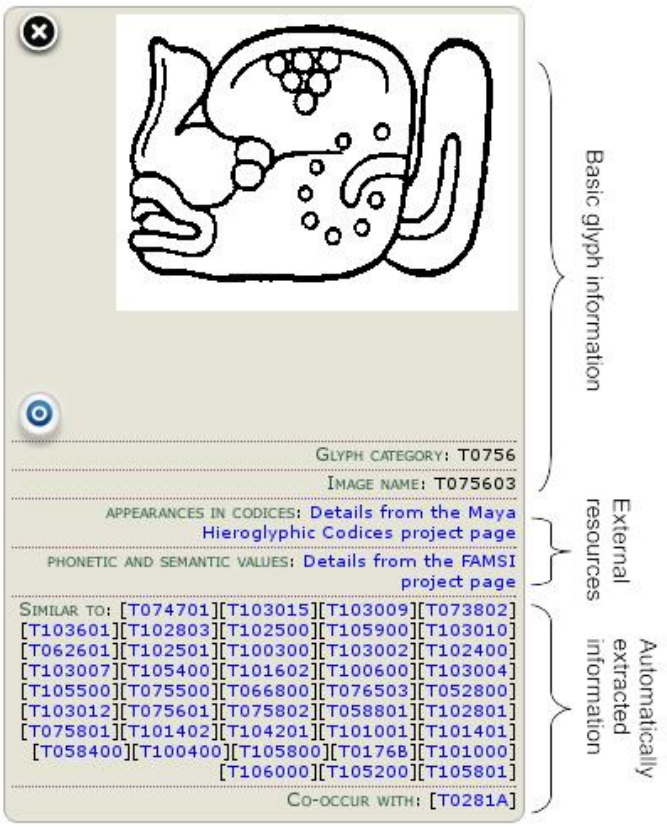

Fig. 11 Pop up information panel of an example glyph, upon clicking on the according node in the graph.

Automatically extracted information.

Additionally, those glyphs that are considered to be visually similar to the current glyph are listed in the pop up information panel (following the 'Similar to' icon). This is automatically recommended based on the methodology presented in section 3.1.1. Glyphs shown from top to bottom, left to right in the list are ordered starting from the most similar one. Clicking on any glyphs in this list will launch an additional pop up information window of the according glyph. Browsing through similar glyphs in the list, users have an opportunity to further explore visually similar glyph patterns.

Last but not least, the statistical glyph co-occurrence information extracted from the three ancient codices is also listed in the information panel (following the 'Co-occur with' icon). The list is ordered from the highest to the lowest frequently co-occurring glyphs to the current glyph. The co-occurrence information provides users an opportunity to explore glyphs in a semantic level by considering the context information. This feature could be particularly useful to help understand problematic glyphs that are difficult to be identified using shape features alone.

3.2 User evaluation

We designed three user evaluation tasks to assess our tool. Our objectives are three-fold. First, to evaluate whether our visualization interface can help users identify unknown glyphs efficiently. Second, to understand whether the automatically recommended visually similar glyphs shown in the information panel indeed share sufficiently similar visual features to the query glyph. Third, to test whether the glyph co-occurrence information provided in the pop up information panel can enable users to identify problematic glyphs by considering the context information within a block. Additionally, we designed questionnaires to collect user feedback to improve our tool.

Two different user groups, namely a general public user group and an epigrapher scholar group, were considered for the evaluation. The general public user group includes participants who have never studied Maya hieroglyphs. For this experiment, we take advantage of two open door exhibition events, one is at Idiap Research Institute and the other at EPFL. These events aim to demonstrate scientific research to the local audience. We set our system in the context of these exhibitions. The audience who participated in these events were from different backgrounds, and their age ranged from around 5 to 70 years old. Once someone from the audience approached our exhibition, we first briefly introduced the Maya culture and hieroglyphs, followed by a general presentation of our project; we then introduce the functionality of our visualization interface. Audiences who were further interested to try our interface, were invited to voluntarily participate in our evaluation task. Within around 8 hours of total exhibition time for each event, there were 26 and 20 volunteers who participated in each study, respectively.

The epigrapher scholar group was formed by experienced Maya epigraphy scholar volunteers. For this user group, we prepared a document which gave detailed explanation of the functionality of our interface, and the designed tasks, followed with a questionnaire. We sent this document via email to scholars possibly interested to test our tool. The evaluation process was carried out by individual scholars independently without any supervision. We collected evaluation results and feedback from three scholars. One of them reported to have less than 3 years' experience, and the other two scholars have 6-10 years´ experience in research on Maya hieroglyphs.

3.2.1 Task I: Identifying unknown glyphs

The first task aimed to assess whether our tool can assist users to identify an unknown glyph efficiently, by searching the catalog glyphs visualized in the graph. Participants were advised to use the shape cluster patterns shown in the graph as a way to reduce the searching space.

Given a query glyph, users first identify a particular shape group in the graph, which this query glyph belongs to based on visual similarity; and then zoom in that particular group to search for the exact match of the query glyph.

For this task, our study targeted both the general public and epigrapher scholars. The evaluation criteria is the amount of time that users spent to find each query glyph.

General public users.

For the general public users, we selected 5 simple query glyphs, each representing a different shape pattern (see Fig. 12). Users were advised to pick any glyph(s) to work with.

From the returned results, we observed that most users tried to search for only one glyph, while some tried to search for more. There were in total 81 attempts to search for individual glyphs by the 46 participants. There was one failed case, in which the user gave up the task after trying for 1.5 minutes. Among the 80 success cases, there were 71 cases where the users were able to find the correct match within 1 minute. The longest time a user took to find a glyph was 3 minutes. We are not aware of any previous study similar to ours in the specific context of Maya hieroglyphs, so this first result could be seen as a future baseline.

Fig. 12 Query glyphs to be identified by general public users.

Epigrapher scholar users.

For the epigrapher scholars, we carefully picked 8 challenging glyphs which can be typically difficult to identify without manually checking the catalog books (see Fig. 13). Epigraphers could pick any glyph(s) from this list to work with. Scholars were advised to record the time they spent to find each query glyph and report their experience.

The evaluation results show that the volunteer scholar who has less than 3 years´ experience was able to find all the 8 glyphs using our tool, among which 6 of them were identified within less than 2 minutes, and the other 2 were identified in around 3 minutes. A second user also reported the time he spent manually searching for the given glyph from the catalog book in comparison to

time spent using our tool. With the tool, this user was able to find 7 glyphs within 1 minute, which is on average 1 minute shorter than his reported time to manually search the catalog. The third user reported his searching experience on one glyph. It took the user 10 seconds to find it with our tool, while he spent more than 5 minutes to find it manually in the catalog. Once again, these times can be seen as a baseline for future comparison.

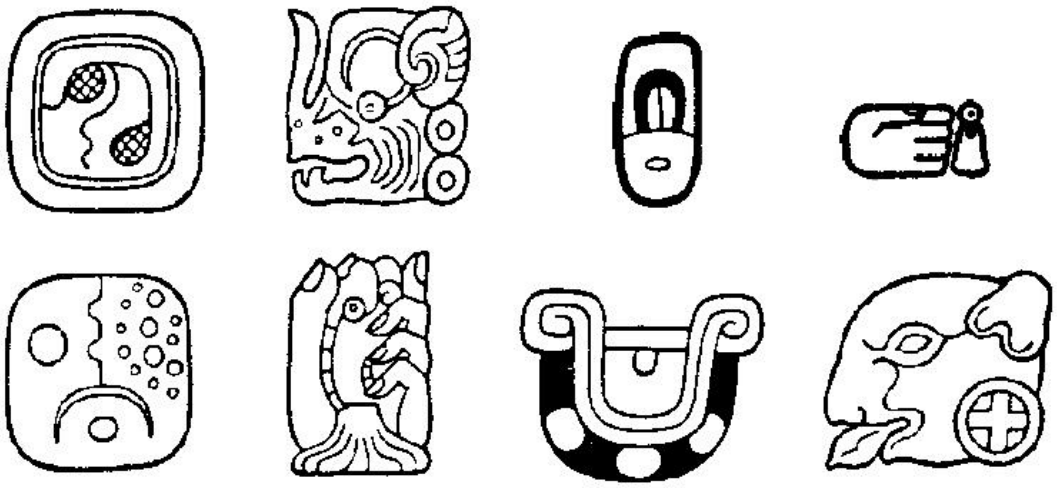

Fig. 13 Query glyphs to be identified by the epigrapher scholar users.

3.2.2 Task II: Evaluating similar glyphs

This task only targeted the epigrapher expert users. We invited participants to manually check each similar glyph listed in the information panel for the 8 glyphs (Fig. 13) which were identified in the previous task, to assess whether they are indeed sufficiently similar to the query glyph. Additionally, participants were also asked to recommend any other visually similar glyphs that were not listed by the tool. On average, there were 25 similar glyphs recommended for each of the 8 query glyphs in our interface. This task aims to further evaluate our methodology to measure the practical value of automatically computed glyph visual similarity.

For each query glyph, the expert participant carefully checked each of the recommended glyphs, marked the ones that were not sufficiently similar to the query, and recommended additional glyphs that were missing in the recommended list. This is a task that experts can perform well. The scholar with less than three years´ experience reported that on average 70% of the recommended glyphs were considered to be visually similar to each of the 8 query glyphs. However, the other two users could only agree on around 20% of the recommended glyphs. The rest were not considered to be sufficiently similar due to various reasons, such as: the recommended glyphs share similar contours to the query, but the contents are not sufficiently

similar (see Fig. 14 for examples); the glyphs share similar shape, but have different visual references (such as a hand shape, a bird shape, etc.). Furthermore, the epigraphers typically recommended 0-3 visually similar glyphs to each of the 8 query glyphs, which were not originally included in the list.

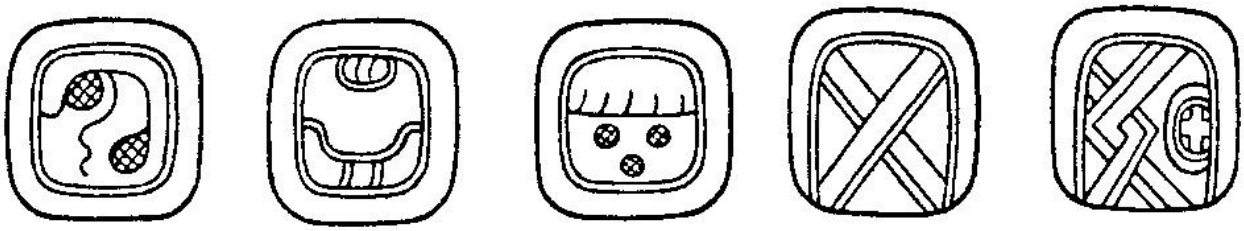

(a) *(b)* *(c)* *(d)* *(e)*

Fig. 14 Example query glyph (a), and negative examples (b-e) from the recommended visually similar glyphs.

3.2.3 Task III: Applying glyph co-occurrence information

This task is designed to evaluate our tool from the perspective of identifying problematic glyphs within the context of glyph blocks, by using the glyph co-occurrence information provided in the information panel. As introduced in section 2.3, context information (glyph co-occurrence) can be used to assist the process of individual glyph identification, especially in cases where the query glyph in a block is difficult to identify due to erosion, occlusion, etc. This task targeted only the epigrapher scholar users.

For this task, we selected 3 partially damaged glyph blocks cropped from ancient Maya codices. The task is to identify each individual glyph in these blocks. Glyph elements in these blocks are ambiguous to be identified due to erosion. The participants were advised to make use of the glyph co-occurrence information provided in the information panel to assist the identification process.

In the results, participants commented that one glyph in one of the blocks was too damaged to be recognized even with the help of the co-occurrence information; in the case of the other 2 blocks, the epigrapher participantss stated that the glyph co-occurrence information provided in our tool was clearly helpful to identify partially damaged glyphs.

3.3 User feedback

We designed questionnaires for the two user groups separately to further evaluate our interface. Our objectives are threefold: first, to learn about the participants' background knowledge with

Maya glyphs; second, to find out whether it was overall a good experience to use our tool; finally, to collect suggestions to improve the tool.

3.3.1 General public user group

From the feedback, 18 users stated that they had previously seen Maya glyphs from videos or images, and had some brief knowledge about the topic; the rest of the users either had never seen any glyphs before, or did not know about Maya writing before this experience. For all users, this was the only tool they had ever encountered to explore and learn about ancient Maya writing. Regarding user experience, users could select one of the five options, on a scale ranging from terrible (1) to great (5). 26 users reported that it had been a great experience, and it had inspired them to learn more about the ancient Maya culture; 18 users claimed that it had been an interesting experience, they have learnt something from it; 2 users reported the experience as being moderately interesting. Regarding the functionality of the tool, 42 users stated that it was a useful tool and they would like to try it again; 4 users commented that it was an interesting tool, but it could be improved. Last but not least, it is worth mentioning a few open comments left by users in hand written form in the questionnaire. One user commented that it would be interesting to apply such interface in museums for visitors to interactively explore the Maya hieroglyphs during their visit. Another user suggested implementing a mobile version of the tool. One participant coincidentally coming from the Yucatan region in Mexico, one of the main Mayan regions, commented that it was a great experience for him to learn about this research project and to explore Maya hieroglyphs through our tool.

3.3.2 Epigrapher scholar user group

From the feedback, our tool was the only one known to all three scholars for exploring Maya hieroglyphs. It was rated as an interesting experience overall, very easy to use, and that it could definitely be applied to assist the daily work of scholars. One participant also commented that the interface could be further improved by incorporating an explicit searching functionality, where users can locate a glyph in the graph by typing the corresponding Thompson glyph category code to the search box.

4. Conclusion

We have introduced the HOOSC descriptor to be used in DH-related shape analysis tasks. We discuss the effect of parameters on the performance of the descriptor in a glyph retrieval framework. Experimental results on ancient Maya hieroglyph data from two different sources

(codex and monument) suggest that a larger spatial context with finer partitioning usually leads to better results, while a smaller spatial context with location information is a good choice for noisy/damaged data.

Additionally, we present a graph-based glyph visualization interface, which enables the exploration and analysis of hieroglyphs. The HOOSC descriptor is used to represent the shape of glyph images, based on which the pairwise glyph similarity scores are computed to define edge weights of the graph.

From the evaluation we conducted with the users, the interface provides an efficient way to browse and explore hieroglyphs. Additionally, Task I and Task II confirmed the effectiveness of the HOOSC descriptor to represent historical Maya glyphs. Results from Task III suggested that the glyph context information provided in our interface can be used to identify problematic glyphs within a block.

The code for HOOSC is available so DH researchers can test the descriptor for their own tasks. Finally, the browsing interface is also publicly available and will be used for further study with the general public and expert users.

Acknowledgement

We thank Edgar Roman-Rangel for co-inventing the HOOSC algorithm and providing the original code. We also thank Guido Krempel, and Jakub Spotak for producing some of the codical data used in this paper. Finally, we thank the Swiss National Science Foundation (SNSF) and the German Research Foundation (DFG) for their support through the MAAYA project.

References

Belongie, S., Malik, J., and Puzicha, J. (2002). Shape Matching and Object Recognition using Shape Contexts. *PAMI.* PP. 509-522.

Bostock M., Ogievetsky V., Heer J. (2011). D3: Data-Driven Documents. *IEEE Trans. Vis. Comput. Graph*. pp. 2301-2309.

Dalal, N., and Triggs, B. (2005). Histogram of Oriented Gradients for Human Detection. *In CVPR*. pp. 886-893.

Eitz, M., Richter, R., Boubekeur, T., Hildebrand, K., and Alexa, M. (2012). Sketch-based shape retrieval. *ACM Transactions on Graphics*. pp. 31:1-31:10

Fischer, A., Bunke, H., Naji, N., Savoy, J., Baechler, M., and Ingold, R. (2012). The HisDoc Project. Automatic Analysis, Recognition, and Retrieval of Handwritten Historical Documents for Digital Libraries. *InterNational and InterDisciplinary Aspects of Scholarly Editing.*

Franken, M., and Gemert, J. C. (2013). Automatic Egyptian Hieroglyph Recognition by Retrieving Images as Texts. *ACM MM*, pp. 765-768.

Hu, R., Odobez, J. M., Gatica-Perez, D. (2016). Assessing a Shape Descriptor for Analysis of Mescoamerican Hieroglyphics: A View Towards Practice in Digital Humanities. *Digital Humanities Conference.*

Hu, R., Can, G., Pallan, C. G., Krempel, G., Spotak, J., Vail, G., Marchand-Maillet, S., Odobez, J. M., and Gatica-Perez, D. (2015). Multimedia Analysis and Access of Ancient Maya Epigraphy: Tools to support scholars on Maya hieroglyphics. *IEEE Signal Processing Magazine*. pp. 75-84.

Montgomery, J., and Helmke, C. (2007). Dictionary of Maya Hieroglyphs, <http://www.famsi.org/mayawriting/dictionary/montgomery/>.

Roman-Rangel, E. (2012). Statistical Shape Descriptors for Ancient Maya Hieroglyphs Analysis. *PhD thesis, EPFL*.

Roman-Rangel, E., Pallan, C. G., Odobez, J. M., and Gatica-Perez, D. (2011). Analyzing Ancient Maya Glyph Collections with Contextual Shape Descriptors. *IJCV*. pp. 101-117.

Stone, A.J. and Zender, M. (2011). Reading Maya Art: A Hieroglyphic Guide to Ancient Maya Painting and Sculpture. *Thames & Hudson Limited Publisher.*

Thompson, J. E. S. (1962). a Catalog of Maya Hieroglyphs. *University of Oklahoma Press*.

Vail, G., and Hernández, C. (2013). The Maya Codices Database, Version 4.1. A website and database available at *http://www.mayacodices.org/*.