

Motif-driven Retrieval of Greek Painted Pottery

S. Lengauer¹, A. Komar¹, A. Labrada³, S. Karl², E. Trinkl², R. Preiner¹, B. Bustos³, T. Schreck¹

¹Institute of Computer Graphics and Knowledge Visualisation, Graz University of Technology

²Institute of Archaeology, University of Graz

³Millenium Institute for Foundational Research on Data, Department of Computer Science, University of Chile

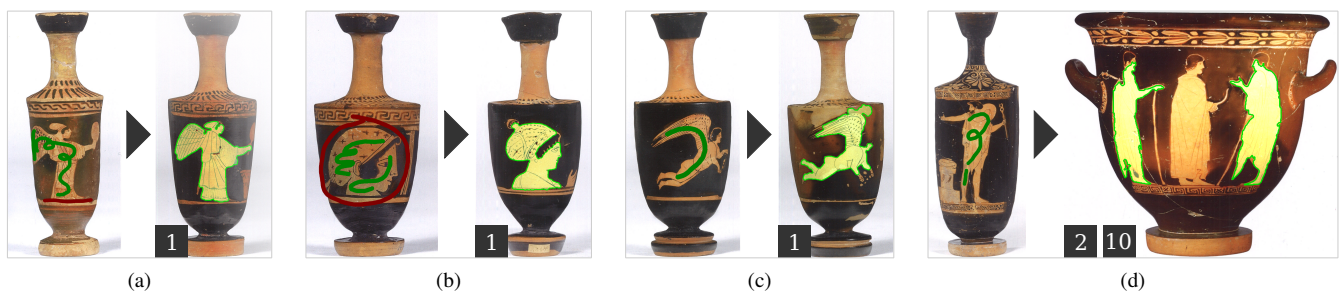


Figure 1: Motif-driven similarity search in ancient pottery databases. (a): A user specifies a query motif of a Nike holding a mirror on a lekythos image via interactive scribbling (left), the top spot of the ranked results identified by our retrieval system (right). Similar examples of query and best match are given for the depiction of a woman's head wearing a sakkos (head scarf) in (b) and the Greek god Eros in (c). Our retrieval system is also capable of finding similar motifs but in mirrored poses by incorporating reflection invariance mechanisms. For the motif of a man with an outstretched arm (d) two motifs with similar poses are detected on the same vessel as second and tenth best match.

Abstract

The analysis of painted pottery is instrumental for understanding ancient Greek society and human behavior of past cultures in Archaeology. A key part of this analysis is the discovery of cross references to establish links and correspondences. However, due to the vast amount of documented images and 3D scans of pottery objects in today's domain repositories, manual search is very time consuming. Computer aided retrieval methods are of increasing importance. Mostly, current retrieval systems for this kind of cultural heritage data only allow to search for pottery of similar vessel's shape. However, in many cases important similarity cues are given by motifs painted on these vessels. We present an interactive retrieval system that makes use of this information to allow for a motif-driven search in cultural heritage repositories. We address the problem of unsupervised motif extraction for preprocessing and the shape-based similarity search for Greek painted pottery. Our experimental evaluation on relevant repository data demonstrates effectiveness of our approach on examples of different motifs of interests.

1. Introduction

The study of Greek painted pottery respectively of Greek vase painting constitutes a major contribution to our understanding of the ancient Greek society [Oak09]. More than one hundred thousand vases are recorded in the Corpus Vasorum Antiquorum (CVA) [cva], providing a world of images allowing us to explore the everyday life as well as the social and religious behavior of this past culture. Since the first studies of Greek vases this subject of figure and narrative art is tackled in various archaeological studies investigating the scenes as well as the figural and ornamental motifs in

chronological, typological, iconographical and mythological terms. Characteristic motifs like mythological figures or draped youths and scenes like wedding or warrior departure, to mention only few, have gained special attention in the Lexicon Iconographicum Mythologiae Classicae (LIMC) [lim] and in many publications (for an overview see [Coo97] or [Boa01]). The basis for all these works is a profound knowledge of materials and to build up repositories of vases with similar motives and scenes. Such repositories allow the application of automatic retrieval system based on exemplary query objects. To date, such retrieval system are mainly driven by

similarity defined by vessel shape, but do not incorporate semantic knowledge of its painted motifs. We present an integrated retrieval system that combines the interactive specification of a query motif by a user with a suitable unsupervised motif extraction pipeline of a search space to allow for a search of specific scenery depicted on painted pottery (see Figure 1). Convolutional Neural networks (CNNs) in general show promising results for these kind of problems (see Section 2). In practice, however the absence of sufficient training data hinders the application of CNN-based approaches for this task. An alternative is to use local features for finding similar keypoint regions across images. On the downside, keypoints are an inappropriate tool for determining the similarity of painted scenes. A more suitable measure of semantic similarity is provided by the similarity of the motif silhouettes. Hence, we follow an approach of segmenting our inputs in a way that each segment corresponds to exactly one individual motif in order to build up a database of motif depictions. By motif, we refer to an ornament or a figure which is itself not a composition or part of another motif. It is crucial that this segmentation step is performed in a robust and reliable automatic way, to be useful for a content-based search engine. In the case of Greek ancient painted pottery, the motifs of interest are mostly painted in two major styles: red figures on a black background (*red-figure pottery*) and vice versa (*black-figure pottery*), yielding objects with a binarized colorization. Hence, a sensible approach is the application of unsupervised gradient-driven segmentation techniques, such as graph-cut based methods as well as segmentation based on morphological transformations. For the retrieval we incorporate color information in addition to shape to rule out background segments with complementary shape of motifs.

In the following, we first discuss related work and then introduce our approach for robust motif segmentation and feature-based retrieval of pottery images. We demonstrate the applicability of our approach by evaluation on a relevant domain image repository, informed by domain experts from Archaeology.

2. Related Work

Our work relates to multimedia retrieval, as well as analytical applications for digital cultural heritage object collections. More specifically, it also relates to approaches for image processing and segmentation. We discuss a selection of previous works in these areas next.

Feature-based Multimedia Retrieval. In content-based multimedia retrieval, a main task is to find multimedia objects satisfying the information need to a user, often expressed by a query. Multimedia retrieval methods help to make use of large amounts of multimedia repositories. A main approach for multimedia retrieval is to represent objects by feature vectors (or descriptors) encoding aspects of interest of the objects, and then comparing objects by differences in their features. Such features can involve low-level measurements extracted from the media (e.g., the color distribution in an image), or higher level features (e.g., concepts within an image). There exist approaches for feature extraction for many different media types, including images [DJLW08], video [LHC*13], or 3D objects [TV08]. Features can be determined locally or globally for a multimedia object. For example, 2D and 3D shape features can be computed for the overall shape, or for local areas on

a shape [NR17, SPS14]. Features can be obtained in unsupervised or supervised ways. In supervised approaches, typically a classifier is trained on a known (labeled) set of objects [LG10]. Then, properties of unknown objects are predicted by the classifier and used as features. In unsupervised approaches, no object labels are required, but the features are computed only from the input data. If appropriate training data is available, supervised features can show to be very effective and often outperform unsupervised approaches [BBZ*17].

Digital Cultural Heritage Data. The digitization of cultural heritage artifacts plays an important role for preservation, presentations, distribution and analysis of objects of interest. To date, image and sketch-based documentation of artifacts is predominant, and many data repositories exist. To name a few, the Corpus Vasorum Antiquorum documents classic potteries for research. Museum institutions are also documenting their contents, sometimes also providing access to the media [bri, met] for wider audiences. Besides images and sketches, 3D digitization is becoming more widespread [BKT13, ham] and opens new possibilities for analysis and comparison of content, e.g., based on geometric analysis approaches [PPY*16]. Similarity-based approaches support many operations on cultural heritage objects, e.g., shape-based search and object association [LKL*19, gra], or object reassembly from 3D fragments [HFG*06, PSA*17]. Besides the geometric shape of objects, also the analysis of painted motifs are of interest for analysis and comparison of cultural heritage objects. Images of motifs from 3D shapes can be obtained by projection onto a 2D viewing plane by appropriate methods [PKBS18]. Using such projections, image-based features can be applied for motif search.

Applications on Archaeological Pottery. Image segmentation and feature recognition of archaeological pottery objects – fragmented or complete – have received more attention in computer science only in the last years. Approaches are mainly based on 2D images, extracting feature vectors [BCT05], visual features of the sherd's surface [PAP*15] and characteristics of the texture information [SBSJ10]. Only recently methods for pattern recognition are extended to 3D models, using an Edge Local Binary Pattern descriptor [TB18, TBS*18]. A general challenge of all these applications in the field of Archaeology is the various preservation stage of the original objects, from perfectly preserved to worn-off surfaces with only small remains of their colorimetric patterns.

Image Segmentation. In image segmentation, the task is to split an image into regions representing meaningful parts. For example, one may wish to isolate a motif (foreground) painted onto a vessel base (background). Many methods for image segmentation have been proposed in image processing [GW06], including construction of region boundaries based on detected edges or growing of regions from seed points based on similarity of texture or other local image properties. Also, morphological operations or fitting of pre-defined generative templates are among the many approaches. Likewise, many techniques exist for 3D object segmentation [CGF09], ranging from simple partitioning schemes [SBS13] to the derivation of 3D skeletons [TDS*16] or graph-based approaches [RKB04, FH04]. In general, there is no universal approach to segmentation, but it has to be guided by application require-

ments, for example, specifying the level of detail at which segments are to be distinguished. Recently, supervised approaches for image segmentation have proven to be successful. For example, in [GDDM14] and [SEZ*14], convolutional neural networks are applied to identify candidate regions of objects within an image. The Convolutional Neural Networks has proven to be an outstanding solution for the semantic segmentation task as well. This task consists of classifying each pixel in an image into a class. Currently, the best approach for semantic segmentation consists in using end-to-end deep neural networks. For example Mask R-CNN [HGDG17] and FPN [LDG*17] are adaptations of end-to-end deep neural networks architecture originally used for object detection. Other important works in the field uses fully convolutional networks architecture (FCN [LSD15], U-NET [RFB15]), this networks contain only convolutional layers trained end-to-end for image segmentation. Finally, a set of very important models for semantic segmentation are all of the deeplab versions, specifically the last one deeplabv3+ [CZP*18] which reported the best results for the 2012 PASCAL VOC segmentation challenge. An implementation of these last models can be found in the open source machine learning library TensorFlow. This library allows to train these models from scratch using own data sets, or to use one of several pre-trained deeplab models for semantic segmentation.

Positioning of this Work. We here present a search engine for motifs on pottery images, for purposes of comparison and retrieval. While our considered data set is not small (hundreds of vessels), obtaining large amounts of labelled training data for supervised approaches is generally difficult. Therefore, we base our approach on an appropriately designed unsupervised workflow for image segmentation, feature extraction, retrieval and result visualization. Our data comprises 2D image documentation from scientific archaeological publications (CVA). The approach is limited by the extent of surface which can be portrayed by a single picture of an object.




Figure 2: Pipeline for generation of our search space of different contour labels. An input image (a) is preprocessed (b) before segmentation (c). (d): Individual segments are extracted and added to our database.

3. Segmentation and Feature Extraction for Pottery Motifs

Our approach poses two major challenges. First, generating the search space by discovering and extracting image segments, possibly corresponding to an individual motif, from domain specific images. Second, the similarity retrieval based on the contour of the extracted segments. The former part encompasses preprocessing and

segmentation (see Figure 2) of a potentially large set of images and therefore needs to be performed purely unsupervised. This poses a particular challenge for the extraction of the motif. At this stage of the preprocessing, classifying image segments as regions of interest and non-interest is not trivial, and also not desirable. Degenerated vessels, bad image conditions, etc. often lead to bad segmentation results which can leave classification attempts cumbersome. Therefore, instead of performing an unsupervised preclassification based on possibly ill-defined classifiers, we include all resulting segments in our search space as possible motifs, thus preserving all the extracted information for the subsequent retrieval.

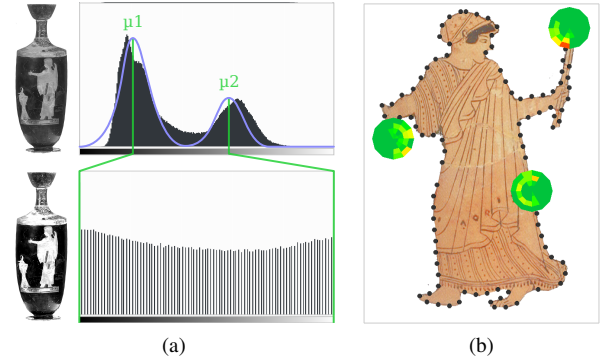


Figure 3: (a): A GMM with two Gaussian components fitted to the original histogram after removal of the background at the top, resulting histogram and image after spreading below. (b) Contour points of a motif with visualization of the shape context descriptors at three locations (green circles). This polar plots indicate the occurrences of other points for specific directions and distances in a heat-map color scheme.

3.1. Preprocessing and Normalization

Our data basis is a series of images depicting pottery with painted surfaces. The motif extraction process needs to be conducted only once for all source images. The aim is to create a database of segments encompassing individual motifs. Together with its position inside the source image, the segment contours are also used for a user friendly result visualization. The aim of the preprocessing is to bring the raw source images into a normalized form so that the parameters of the subsequent segmentation can be trimmed to a generalized input. In detail, this involves scaling, removal of noise, enhancement of important and discardment of unimportant information.

Firstly, the images are scaled to uniform height of 512 pixels which was established to be an appropriate trade-off between efficiency and preservation of low-level details. After converting the image to gray scale the primal-dual denoising algorithm, presented by [CP11], is applied in order to remove noise while preserving essential image features. Subsequently, the contrast between motifs and background is maximized. As proposed by Otsu [Ots79] we assume the presence of two dominant color classes corresponding to the mean color of the vessel surface and the mean color of the motif paintings, referred to by μ_1 and μ_2 respectively in the




Figure 4: Examples of our inputs (top) with the segmentation results for EGBIS (middle) and the morphological segmentation (bottom).

histogram of one particular example given in Figure 3a. To obtain the maximal contrast between motif and background the original spectrum is spread between these two points. The position of the maxima is obtained by fitting a Gaussian Mixture Model (GMM) with means μ_k and covariances Σ_k with 2 Gaussian components to the grayscale color histogram of the depicted vessel. The positions of μ_1 and μ_2 are obtained with an expectation maximization approach where the most likely model parameters of the GMM $\theta_{\text{ML}} = \{\alpha_1, \alpha_2, \mu_1, \mu_2, \Sigma_1, \Sigma_2\}$ are estimated by maximizing the logarithmic probability $\ln P(\mathcal{X}|\theta)$ of observing the histogram distribution $\mathcal{X} = \{x_1, \dots, x_N\}$ given the parameters θ :

$$\theta_{\text{ML}} = \arg \max_{\theta} \{\ln P(\mathcal{X}|\theta)\}. \quad (1)$$

μ_1 and μ_2 are taken as the new min and max values of the histogram. All values outside the interval $[\mu_1, \mu_2]$ are set to μ_1 or μ_2 respectively and all values inside are interpolated linearly. The result of this step is depicted in the bottom row of Figure 3a.

3.2. Motif Segmentation

A central task for the motif extraction preprocessing is the segmentation of the input images, ideally in such a way that one segment corresponds to exactly one motif. In general, a suitable image segmentation poses a major challenge in the field of computer vision with a variety of approaches, all exhibiting different data-dependent strengths and weaknesses. In the context of our work we apply and evaluate two different segmentation techniques: a graph-based approach utilizing gradient information of the grayscale image and an approach based on morphological transformations working on the binarized input.

Graph-cut based Segmentation. The idea of graph-based approaches is to treat an image as an unconnected graph $G = (V, E)$ where the vertices V correspond to pixels and the edges E are used to model the neighborhood of pixels. The edges carry a weight which is a measure for the dissimilarity between two pixels. The

aim of the segmentation is to split the graph into a series of connected components where the elements of each component are as similar as possible. Our implementation of a graph-cut based segmentation (EGBIS) algorithm is based on the Efficient Graph-Based Image Segmentation (EGBIS) algorithm described by [FH04]. The behavior of the unsupervised segmentation can be governed by adjusting the minimum size of a component $|C|$ (in number of pixels) and the dimensionless scaling parameter k which models a tendency to larger components. We adjusted these parameters empirically for our set of inputs consisting of about 100 images depicting various vessel types exhibiting different motifs with some examples given in the top row of Figure 4. A good segmentation for a majority of the motifs (see second row of Figure 2e) was obtained with $|C| = 1500$ and $k = 1500$. The segment corresponding to the scene background behind the displayed object is determined to be the segment containing the four corner pixels of the input image and is consequently discarded. All other resulting segments are subsequently treated individually by taking the segmented area as foreground mask for a foreground/background segmentation, described by [RKB04]. Using the input image in full resolution and scaling the mask accordingly allows us to preserve fine details. The segments resulting from this step are added to our database.

Morphological Segmentation. Although the graph-cut based performs well in general, it is specifically sensitive to fine gaps, i.e., lines in the gradient image, which can pose a problem for specific pottery art exhibiting fine stylistic strokes. As shown in Figure 5a, this can lead to an over-segmentation, separating the depicted Eros motif from the belonging wings. In the specific case of ancient Greek pottery, the motifs of interest are mostly depicted in dual-colored images, exhibiting either black figures on the red clay background or vice versa. This kind of appearance is specifically suitable for a reduction to binary images, which allows for a straight-forward extraction of connected foreground components. Directly labeling the binary image of the normalized grayscale output from Section 3.1 already shows a better preservation of the figure's shape, but has similar problems in the face of gaps, pro-

ducing a more complex silhouette as seen at the stylized feathers (Figure 5b). To this end, we apply a series of morphological operations to the initial binary image. In the first step, we downscale the images to half the resolution for efficiency, and apply a closing operation using a small disk-shaped structure element (2 pixel radius) to fill out fine recurring gaps in the motifs. As shown in Figure 5c, this reattaches the leg of the figure, and closes the silhouette of the wings. In the second step, we add a morphological opening using a larger kernel (6 pixel radius). This has two effects: First, it produces a safer detachment of individual motifs. Second, it smooths the silhouette and removes fine high frequency details like the Kithara element held by Eros in Figure 5d, which is beneficial for the retrieval described in the following section. Exemplary results for this segmentation technique are given in the bottom row of Figure 4. Similar to the EGBIS approach small residual segments below a minimum threshold size (1% of the image area in all examples) are discarded.




Figure 5: Comparison of different segmentation approaches for an Eros motif. (a) Graph-cut based segmentation. (b) Connected-component labeling of the binarized input image, (c) after small-scale closing, (d) after additional large-scale opening.

3.3. Similarity Retrieval of Segmented Motifs

We found that the shape of their contours is a simple and robust property for determining the similarity of motifs. Thus, we base our retrieval mostly on the contour outlines of motifs. As a result of the above segmentation, we have a reference database of motifs for our experiments.

Query Specification. In our prototype the query input by the user consists of a single image containing the motif to search for. To specify the image regions containing the motif we provide the user with an interface which allows to select regions belonging to the motif as well as regions containing irrelevant scenery by means of a selection brush with adjustable brush size (see left side of Figure 6a). In most cases a very coarse selection with a few brush strokes produces sufficiently good results. The selections serve as foreground/background mask for a *GrabCut* foreground extraction described by [RKB04]. The interface shows the preliminary results of this step in a separate view (see right side of Figure 6a) In cases where the proposed segmentation does not match the user’s expectations, the selection can be refined iteratively by adding additional brush strokes. In Figure 6b it can be seen that the hair of the depicted character has approximately the same saturation as the remaining vessel surface, leading to an erroneous segmentation with

the initial coarse selection (Figure 6b, upper row). By explicitly marking those areas as foreground, an improved selection can be obtained as depicted in the bottom row of Figure 6b. The outline of the foreground of this segmentation is taken as the query for the subsequent retrieval. Alternative to this input module, motifs depictions which have been extracted by other means (e.g. graphics editors) could be used for retrieval.




Figure 6: (a): A coarse selection of foreground or motif areas (green) and background areas (red) with the resulting segmentation (background marked in green). (b): An erroneous segmentation (top row) can be further refined by iteratively specifying regions (bottom row).

Feature Extraction. For similarity matching we rely on shape as well as color. In terms of shape the *shape context* feature descriptor, described by [BMP02], has been found to be a good choice for this application. This local shape feature descriptor takes a representation of a shape by a number of points and defines a feature vector for each point based on the relative location, directivity and distance of all the other points in form of a distribution histogram with b_r radial bins and b_θ angular bins. An illustration of this descriptor is given in Figure 3b. In contrast to *keypoints* (“significant” points in an image) which are used by most local feature descriptors, the points used for the shape contexts are taken uniformly along the contour. In our case, only the external contour Ω of the shape has been used with a fixed number of points of $N = 100$ which is a trade-off between accuracy and computational performance. To this end we took every n -th point of the external contour, where

$$n = \left\lceil \frac{|\Omega|}{N} \right\rceil \quad (2)$$

with $\lceil \cdot \rceil$ denoting the round operation and $|\Omega|$ the number of contour points. If the resulting number of points is too high, randomly selected points are removed until the required number is met. If the resulting N is too low, points are added iteratively by randomly selecting a point and placing a new point at a linearly interpolated position between him and his neighbor. In terms of histogram resolution of the feature descriptor $b_r = 4$ and $b_\theta = 12$ have been chosen for our evaluation.

Reflection Invariant Similarity Score. With a descriptor assigned to each point of a query q_i and each point t_j of a specific target, the cost of matching the points $C_{ij} = C(q_i, t_j)$ can be estimated. Since histograms are compared, the χ^2 distance is an appropriate metric. The total cost for assigning all points in the query to

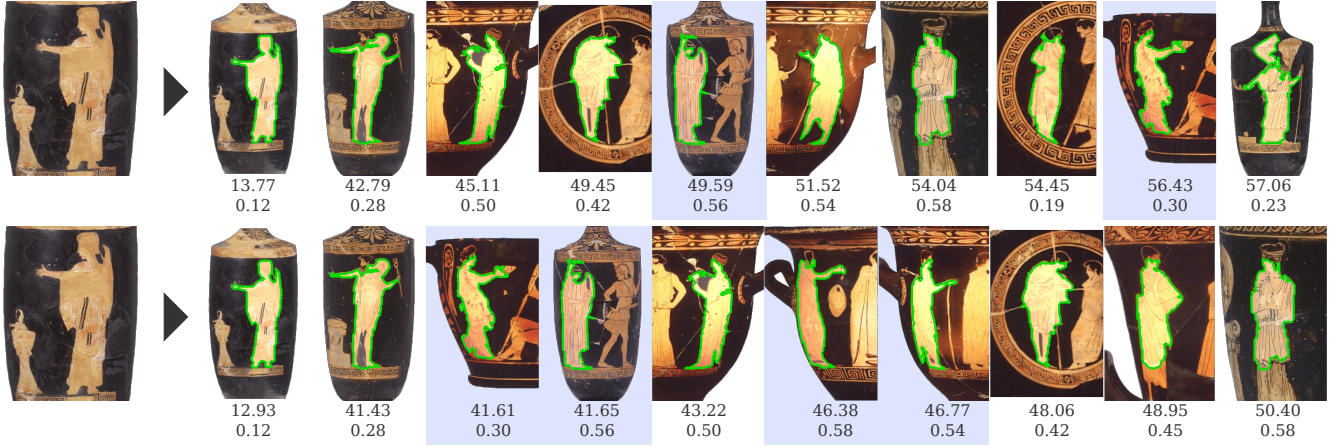


Figure 7: Query motif of a person with outstretched arm with retrieval results and distance metrics for the baseline shape context descriptor and Bhattacharyya distance (top row) and with adjustment for reflection invariance (bottom row). Our approach is able to identify relevant motifs from the target database, including reflected motifs. Blue highlights indicate figures stretching their arm in the opposite direction of the query figure.

all points in the target with assignment π is then given by the sum of the individual costs

$$\tilde{H}(\pi) = \sum_i C(q_i, t_{\pi(i)}) \quad (3)$$

over all points i , which is an inverse measure of similarity. Finding the optimal assignment (the assignment with the lowest costs) is a square assignment problem. It can be solved with the Hungarian algorithm [Kuh55] with a runtime complexity of $\mathcal{O}(n^3)$. For our experiments we applied the more efficient method proposed by Jonker-Volgenant [JV87] which is also of cubic complexity but drastically reduces computational load through several optimizations. The cost of the optimal assignment was used as a metric for ranking different results.

Since we did no filtering when adding segments to the database there are also a lot of segments with the vessel background in-between or around motifs. Those can share major parts of their contour line with a motif (complementary contour) and are thus likely to also provide a good match in terms of shape while being an undesired result. At the same time, in such cases there is a huge difference in color distribution between query and target segment. Thus, we combine our shape-based retrieval with a color based approach to discriminate complementary segments. From the 256 bin histograms of gray scale query- and target segment the Bhattacharyya distance [Bha43] was computed. Target segments having a distance of equal to or more than 0.6 to the query segment are excluded from the result set.

From inspection of the target search space and discussion with domain experts it was concluded that motifs often appear in similar poses but horizontally flipped. Figure 1d gives an example, specifically, a figure with outstretched arm facing leftwards and another one facing rightwards. We made our retrieval invariant towards this kind of transformation by considering a flipped version of our query. This can be efficiently done by reversing the order of the angular bins b_θ of our query descriptors, giving an alternative

sum of costs $\tilde{H}_{mirror}(\pi)$. The overall cost with respect to a specific target is then given by

$$H(\pi) = \min\{\tilde{H}(\pi), \tilde{H}_{mirror}(\pi)\}. \quad (4)$$

4. Experimental Evaluation

We have selected two recent published CVA volumes regarding red-figured vases, exhibiting a representative range of various motifs on different vessel's shapes, like lekythoi as well as open shapes like kraters, cups, plates or pyxides. Those publications are the Corpus Vasorum Antiquorum Berlin 13 [ZE13] and Corpus Vasorum Antiquorum Dresden 3 [Esc18]. Starting from page scans containing multiple depictions and meta information, individual images have been extracted as described in [LKL*19]. Almost all of the depicted vessels exhibit motifs on their surfaces. All images depicting whole vessels with motifs were manually selected, resulting in a data basis of 57 images from *CVA Berlin 13* and 42 images from *CVA Dresden 3*. A total of 152 individual image segments are extracted with morphological segmentation and 785 with EGBIS. This discrepancy is due to higher degree of fragmentation and due to the fact that segments containing vessel background are present with the latter segmentation method. We evaluate different aspects of our retrieval system using different query shapes, including both figural and ornamental query motifs to evaluate. We evaluate the aspects of similarity scoring as well as different segmentation methods of our retrieval system using both figural and ornamental query motifs.

Figural Query Motifs. As a first query example we choose a depiction of a person exhibiting a characteristic pose that often appears in our domain image space. Figure 7 (left) illustrates the query image of a human figure with an outstretched arm to the left, a gesture of speech. The right side shows the ten best ranked retrieval results in descending order with the found segments outlined




Figure 8: Retrieval results for a palmette leaf appearing on opposite sides of the same vessel. Glyphs above the first seven result items indicate front or back view of the depicted vessel.

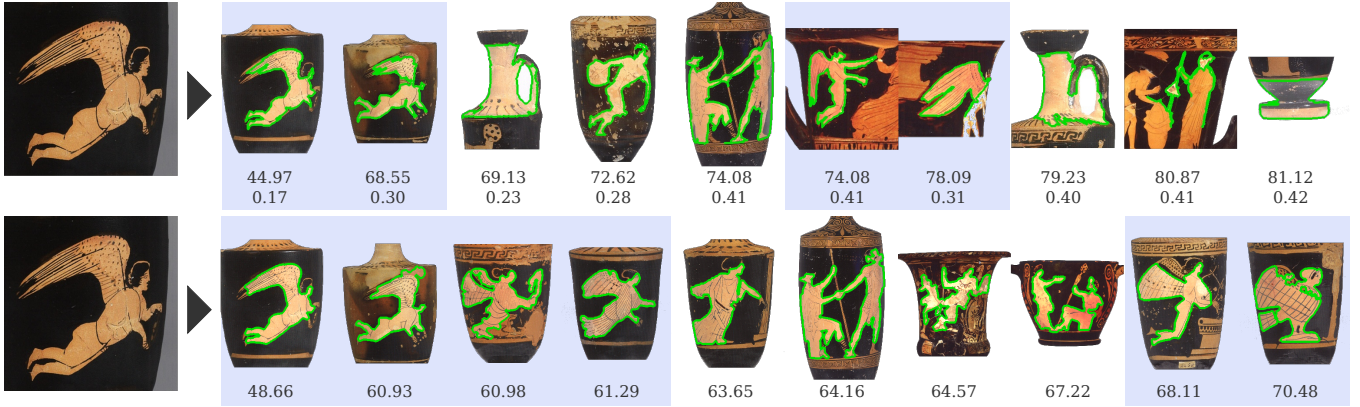


Figure 9: Retrieval results for a depiction of Eros as query over a search space segmented with the EGBIS technique (top row) and a search space segmented with morphological operations. A blue highlight indicates the presence of a motif depicting a winged figure.

in green. The numbers below the figures denote the similarity score (i.e., the assignment costs) $H(\pi)$ defined in Equation (3) as well as the Bhattacharyya distance measuring the histogram similarity of the segments. The upper row shows the results using the baseline shape contexts descriptor (see Section 3.3).

As expected, the top ranked result depicts the query object itself. Overall, the top ten results all correspond to figural motifs, with several figures exhibiting variations relevant to the query motif, as confirmed by the domain research co-authors. Due to the reflection variant nature of the baseline feature descriptor, the found objects mostly depict figures outreaching their arm to the same direction.

Reflection Invariance. When extending the similarity score to the reflection invariant form in Equation (4) we obtain a ranked result shown in the lower row of Figure 7. A quantitative improvement of the result set can be immediately read from the tighter distribution of top similarity scores over a lower range. Qualitatively this is due to new high ranked objects (highlighted in blue) that particularly depict figures with an arm outreaching to the *right*. Another example is given by an alternative query for a palmette leaf (see Figure 8). This particular ornament type appears in multiple depictions in the search space covering two palmettes on opposite sites of the same vessel, mostly below the handles of the vessel. In this specific case the query finds multiple similarities on both sides of the vessel, and – due our reflection invariant descriptor – mirror symmetric correspondences within each palmette. This example especially demonstrates the capability of our similarity search for other

applications, such as unsupervised object based clustering of detail images.

Evaluation of Segmentation Methods. As a particularly interesting query example we chose an often reoccurring mythological figure on Greek pottery, which is the depiction of Eros, the Greek god of love, appearing as a young male figure with wings. This motif exhibits a more complex shape silhouette which is particularly difficult to capture both by shape descriptors and segmentation methods. We use this example for a comparative evaluation of the two segmentation approaches we consider in our approach (see Section 3.2). The first row in Figure 9 shows the results for this motif using the EGBIS technique. After the query motif itself on rank one, this segmentation leads to a result set containing three further Erotes in the top ten ranking (blue highlights). Within this top ten, we observe results depicting complex motifs that do not directly resemble a characteristic winged figure, which cannot be sufficiently discriminated by our shape descriptor. Moreover we also find three clearly unrelated image segments corresponding to the mouths and the feet of different vessels, which exhibit a similar shape as well. This is a result of our design choice to include *all* found segments in the search space. In the second row of Figure 9 the same query is run on a morphologically segmented database. Compared to the EGBIS based result we now retrieve two Erotes, two depictions of a Nike (winged goddess of victory) and a depiction of a Siren (mythological hybrid of woman and bird) after the top ranked query object. On rank 9 is an Eros image exhibit-

ing a high similarity score that has not been present in the previous top ten result set. As discussed in Figure 5 this is due to an over-segmenting of its wings, and therefore changing the overall shape. Our shape descriptor finds hit 3 and 10 even in the face of partially degenerated segmentations. As before, we find seemingly similar motifs of complex shapes, but no unrelated mouth segment within the top ten. Overall, the morphology-based segmentation proves to perform better for this query, which is also reflected in the tighter similarity distance distribution within the top ten list.

5. Limitations

Motifs rarely occur unattached to the surrounding scenery but are oftentimes either connected to other motifs or to other (more abstract) decorative paintings as given in Figure 10a. This poses a severe problem for the graph-cut segmentation approach which will in many cases set no cuts in those border case regions. With ancient pottery a common issue is that the surfaces exhibit characteristic degenerations in the form of cracks and split-offs (see Figure 10b). This means parts of a motif are possibly missing and at the same time artificial gradients are introduced which are again a difficulty for segmentation.




Figure 10: (a) Two motifs intersecting each other (green circle) and parts of other decorative paintings (blue circle). (b) Crack intersecting with motif (red circles).

Our applied shape contexts feature descriptor is variant to non-rigid transformations. Hence, similar figures but in different poses are usually not detected. High frequency details of a query cannot be represented appropriately since the number of contour points N needs to be sufficiently low for an efficient retrieval. The bottleneck in terms of runtime is the high complexity of the square assignment. While for $N = 100$, feature extraction from the segmented contours takes about 7.5 ms per motif, and square assignment about 25.1 ms, $N = 150$ already requires 16.6 ms and 94.7 ms, respectively.

The pottery painting images we use as an input are distorted due to the curved surface of the vessels. On the one hand this introduces unwanted transformations to the motifs, on the other hand motifs can usually not be well captured with a single photograph but multiple images from different views are necessary (see Figure 11). At the current state the belonging of multiple images to one and the same object is not modelled in our search space. It is




Figure 11: Multiple pictures for one and the same motif of a woman with sakkos and torch from different viewpoints.

thus possible (however unlikely) to get multiple matches for the same motif on the same vessel but from different perspectives in the retrieval result set. Part of our future work involves the transfer from an image-based to an object-based search space. This allows us to discard a overhead for matches for one and the same motif.

6. Future Work

To overcome the limitations of surface distortion and truncated motif depictions we want to expand our search space with drawings of pottery paintings which are complete and stylized version of the motif [BKMM18]. Alternatively, a rollout or flattening of the original 3D object, as described by [PKBS18], could be used as input. The part of our approach having the highest potential for improvement is the segmentation. Both our used methods come with different strengths and weaknesses. In addition to combining them in a way that utilizes their individual advantages, we also want to look into other methods, especially such based on deep neural networks. To cope with unwanted gradients introduced by cracks or split-offs the three dimensional information of high resolution models could be exploited. Furthermore, the usage of specialized features to describe and possibly help segment ornamenting background could be explored. For example, in [?] an approach to describe and retrieve ornament shapes in 3D data is proposed. A similar approach could be devised and taken into account for an improved version of our motif-based search. We also want to address the non-rigid transformation invariance of our applied feature descriptor such that it can cope with this kind of transformation as proposed by [GAW15]. To this end, other features altogether or combinations of those could be used. In the process we also want to look into ways to incorporate the edges and texture information inside a contour. Other interesting future work includes approaches for visual cluster analysis for over-viewing large amounts of motifs, and comparative shape visualization techniques for in detail comparison of resulting motifs.

7. Conclusion

We have presented an interactive retrieval system that makes use of the motif information embedded in the surface of Greek painted pottery. The two major challenges to this end contain the unsupervised extraction and segmentation of domain relevant image data to build a database of motif segments as well as the task of conducting a shape-based retrieval on such a database. For the task of segmentation we have evaluated two different methods, including a graph-cut based as well as a morphological segmentation approach. For the retrieval, off-the-shelf color and shape feature descriptors have been adjusted to our needs. Our experiments show that the resulting retrieval system finds semantically relevant images for different query motifs, on a representative target search space. We also discussed limitations and possible extensions of our approach.

Acknowledgments

This work was co-funded by the Austrian Science Fund FWF and the State of Styria, Austria within the project *Crossmodal Search and Visual Exploration of 3D Cultural Heritage Objects* (P31317-NBL) as well as the Millennium Institute for Foundational Research on Data (IMFD), Chile.

References

- [BBZ*17] BAI S., BAI X., ZHOU Z., ZHANG Z., TIAN Q., LATECKI L. J.: GIFT: Towards scalable 3D shape retrieval. *IEEE Trans. Multimedia* 19, 6 (2017), 1257–1271. doi:10.1109/TMM.2017.2652071. 2
- [BCT05] BISHOP G., CHA S.-H., TAPPERT C.: A greek pottery shape and school identification and classification system using image retrieval techniques. *Proceedings of Student/Faculty Research Day CSIS. Pace University* (2005). 2
- [Bha43] BHATTACHARYA A.: On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math. Soc.* 35 (1943), 99–109. 6
- [BKMM18] BAYER P. V., KARL S., MARA H., MÁRTON A.: Advanced documentation methods in studying corinthian black-figure vase painting. <https://archiv.ub.uni-heidelberg.de/volltextserver/25189/>, 2018. Accessed: 2019-07-19. 8
- [BKT13] BREUCKMANN B., KARL S., TRINKL E.: Digitising ancient pottery. Precision in 3D. *Forum Archaeologiae* 66, III (2013). URL: <http://farch.net>. 2
- [BMP02] BELONGIE S., MALIK J., PUZICHA J.: Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 4 (2002), 509–522. 5
- [Boa01] BOARDMAN J.: *The history of Greek vases: potters, painters, and pictures*. Thames & Hudson, 2001. 1
- [bri] British Museum online research collection. https://www.britishmuseum.org/research/collection_online/search.aspx. Accessed: 2017-09-04. 2
- [CGF09] CHEN X., GOLOVINSKIY A., FUNKHOUSER T.: A benchmark for 3D mesh segmentation. *ACM Transactions on Graphics (Proc. SIG-GRAPH)* 28, 3 (Aug. 2009). 2
- [Coo97] COOK R. M.: *Greek painted pottery*. Routledge, 1997. 1
- [CP11] CHAMBOLLE A., POCK T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision* 40, 1 (2011), 120–145. 3
- [cva] Corpus Vasorum Antiquorum. <http://www.cvaonline.org/cva/>. Accessed: 2019-07-07. 1
- [CZP*18] CHEN L.-C., ZHU Y., PAPANDREOU G., SCHROFF F., ADAM H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)* (2018), pp. 801–818. 3
- [DJLW08] DATTA R., JOSHI D., LI J., WANG J.: Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys* 40, 2 (2008), 1–60. 2
- [Esc18] ESCHBACH N.: *CVA Dresden 3: Attisch rotfigurige Keramik*. C. H. Beck, 2018. 6
- [FH04] FELZENSZWALB P. F., HUTTENLOCHER D. P.: Efficient graph-based image segmentation. *International Journal of Computer Vision* 59, 2 (Sep 2004), 167–181. doi:10.1023/B:VISI.0000022288.19776.77. 2, 4
- [GAJW15] GUERRERO P., AUZINGER T., WIMMER M., JESCHKE S.: Partial shape matching using transformation parameter similarity. *Computer Graphics Forum* 34, 1 (2015), 239–252. doi:10.1111/cgff.12509. 8
- [GDDM14] GIRSHICK R., DONAHUE J., DARRELL T., MALIK J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition* (Washington, DC, USA, 2014), CVPR '14, IEEE Computer Society, pp. 580–587. doi:10.1109/CVPR.2014.81. 3
- [gra] GRAVITATE EU project on reconstruction and re-unification of shattered cultural heritage objects. <http://gravitate-project.eu/>. Accessed: 2017-08-04. 2
- [GW06] GONZALEZ R. C., WOODS R. E.: *Digital Image Processing (3rd Edition)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2006. 2
- [ham] Virtual Hampson Museum. <http://hampson.cast.uark.edu/>. Accessed: 2017-08-04. 2
- [HFG*06] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling fractured objects by geometric matching. *ACM Trans. Graph.* 25, 3 (July 2006), 569–578. doi:10.1145/1141911.1141925. 2
- [HGDG17] HE K., GKIOXARI G., DOLLÁR P., GIRSHICK R.: Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision* (2017), pp. 2961–2969. 3
- [JV87] JONKER R., VOLGENANT A.: A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing* 38, 4 (1987), 325–340. 6
- [Kuh55] KUHN H. W.: The hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1-2 (1955), 83–97. doi:10.1002/nav.3800020109. 6
- [LDG*17] LIN T.-Y., DOLLÁR P., GIRSHICK R., HE K., HARIHARAN B., BELONGIE S.: Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 2117–2125. 3
- [LG10] LIM K., GALOOGAHI H. K.: Shape classification using local and global features. In *2010 Fourth Pacific-Rim Symposium on Image and Video Technology* (2010), pp. 115–120. doi:10.1109/PSIVT.2010.26. 2
- [LHC*13] LIU J., HUANG Z., CAI H., SHEN H. T., NGO C. W., WANG W.: Near-duplicate video retrieval: Current research and future trends. *ACM Computing Surveys* 45, 4 (Aug. 2013), 44:1–44:23. doi:10.1145/2501654.2501658. 2
- [lim] Lexicon Iconographicum Mythologiae Classicae. <http://www.limc.ch>. Accessed: 2019-07-07. 1
- [LKL*19] LENGAUER S., KOMAR A., LABRADA A., KARL S., TRINKL E., PREINER R., BUSTOS B., SCHRECK T.: Sketch-Aided Retrieval of Incomplete 3D Cultural Heritage Objects. In *Eurographics Workshop on 3D Object Retrieval* (2019), Biasotti S., Lavouaf G., Veltkamp R., (Eds.), The Eurographics Association. doi:10.2312/3dor.20191057. 2, 6

- [LSD15] LONG J., SHEHAMER E., DARRELL T.: Fully convolutional networks for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2015). 3
- [met] Metropolitan Museum of Art online collection. <http://www.metmuseum.org/art/collection>. Accessed: 2017-09-04. 2
- [NR17] NEAL F. B., RUSS J. C.: *Measuring Shape*. CRC Press, 2017. 2
- [Oak09] OAKLEY J. H.: Greek vase painting. *American Journal of Archaeology* 113, 4 (2009), 599–627. 1
- [Ots79] OTSU N.: A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics* 9, 1 (Jan 1979), 62–66. doi:10.1109/TSMC.1979.4310076. 3
- [PAP*15] PICCOLI C., APARAJEYA P., PAPADOPOULOS G. T., BINTLIFF J., LEYMARIE F. F., BES P., VAN DER ENDEN M., POBLOME J., DARAS P.: Towards the automatic classification of pottery sherds: two complementary approaches. *Across Space and Time* (2015), 463. 2
- [PKBS18] PREINER R., KARL S., BAYER P., SCHRECK T.: Elastic flattening of painted pottery surfaces. In *Proc. EG Workshop on Graphics and Cultural Heritage* (2018). Short paper. 2, 8
- [PPY*16] PINTUS R., PAL K., YANG Y., WEYRICH T., GOBBETTI E., RUSHMEIER H.: A survey of geometric analysis in cultural heritage. *Computer Graphics Forum* 35, 1 (2016), 4–31. doi:10.1111/cgf.12668. 2
- [PSA*17] PAPAIOANNOU G., SCHRECK T., ANDREADIS A., MAVRIDIS P., GREGOR R., SIPIRAN I., VARDIS K.: From reassembly to object completion - a complete systems pipeline. *Computing and Cultural Heritage* 10, 2 (February 2017). 2
- [RFB15] RONNEBERGER O., FISCHER P., BROX T.: U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (2015), Springer, pp. 234–241. 3
- [RKB04] ROTHER C., KOLMOGOROV V., BLAKE A.: "grabcut": Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 309–314. doi:10.1145/1015706.1015720. 2, 4, 5
- [SBS13] SIPIRAN I., BUSTOS B., SCHRECK T.: Data-aware 3D partitioning for generic shape retrieval. *Computers & Graphics Special Issue on 3D Object Retrieval* 37, 5 (August 2013), 460–472. doi:10.1016/j.cag.2013.04.002. 2
- [SBSJ10] SMITH P., BESPALOV D., SHOKOUFANDEH A., JEPPSON P.: Classification of archaeological ceramic fragments using texture and color descriptors. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops* (2010), IEEE, pp. 49–54. 2
- [SEZ*14] SERMANET P., EIGEN D., ZHANG X., MATHIEU M., FERGUS R., LECUN Y.: Overfeat: Integrated recognition, localization and detection using convolutional networks. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings* (2014). URL: <http://arxiv.org/abs/1312.6229>. 3
- [SPS14] SAVELONAS M., PRATIKAKIS I., SFIKAS K.: An overview of partial 3D object retrieval methodologies. *Multimedia Tools and Applications* (2014), 1–26. 2
- [TB18] THOMPSON E. M., BIASOTTI S.: Edge-based lbp description of surfaces with colorimetric patterns. *arXiv preprint arXiv:1804.03977* (2018). 2
- [TBS*18] THOMPSON E. M., BIASOTTI S., SORRENTINO G., POLIG M., HERMON S.: Towards an Automatic 3D Patterns Classification: the GRAVITATE Use Case. In *Eurographics Workshop on Graphics and Cultural Heritage* (2018), Sablatnig R., Wimmer M., (Eds.), The Eurographics Association. doi:10.2312/gch.20181372. 2
- [TDS*16] TAGLIASACCHI A., DELAME T., SPAGNUOLO M., AMENTA N., TELEA A.: 3D Skeletons: A State-of-the-Art Report. *Computer Graphics Forum* (2016). doi:10.1111/cgf.12865. 2
- [TV08] TANGELDER J., VELTKAMP R.: A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications* 39, 3 (2008), 441–471. doi:10.1007/s11042-007-0181-0. 2
- [ZE13] ZIMMERMANN-ELSEIFY N.: *CVA Berlin 13: Attisch rotfigurige Lekythen*. C. H. Beck, 2013. 6