

Graph-Based Shape Similarity of Petroglyphs

Markus Seidl¹, Ewald Wieser¹, Matthias Zeppelzauer¹,
Axel Pinz², and Christian Breiteneder³

¹ St. Pölten University of Applied Sciences, St Pölten, Austria
`markus.seidl@fhstp.ac.at`

² Graz University of Technology, Graz, Austria

³ Vienna University of Technology, Vienna, Austria

Abstract. Petroglyphs can be found on rock panels all over the world. The possibilities of digital photography and more recently various 3D scanning methods opened a new stage for the documentation and analysis of petroglyphs. The existing work on petroglyph shape similarity has largely avoided the questions of articulation, merged petroglyphs and potentially missing parts of petroglyphs. We aim at contributing to close this gap by applying a novel petroglyph shape descriptor based on the skeletal graph. Our contribution is twofold: First, we provide a real-world dataset of petroglyph shapes. Second, we propose a graph-based shape descriptor for petroglyphs. Comprehensive evaluations show, that the combination of the proposed descriptor with existing ones improves the performance in petroglyph shape similarity modeling.

Keywords: Petroglyph similarity · Shape similarity · Graph matching · Graph edit distance · Graph embedding

1 Introduction

Petroglyphs have been pecked, scratched and carved in rock panels all over the world. The documentation with digital photography and more recently various 3D scanning methods enabled (semi-)automated analysis of petroglyphs and attracts research activity. In the past few years, works related to segmentation of petroglyph images [29], automated classification of petroglyph shapes [38][39][9][10] and digital presentation of rock art [30] have been published. The large number of single petroglyphs makes the usage of automated analysis methods attractive. In this paper, we contribute to the evaluation of shape similarity of petroglyphs. Our contribution is twofold: First, we provide a real-world dataset of petroglyph shapes that have been digitized from tracings of petroglyphs of the UNESCO world heritage site Valcamonica. The dataset further contains fine-grained expert annotations from archeologists. Second, we propose and evaluate a graph-matching approach for shape similarity of petroglyphs.

The motivation to investigate the skeletal graphs of petroglyphs for shape matching is with respect to our ultimate goal: The retrieval of possibly merged and unfinished or damaged (i.e. partial) petroglyphs from full rock panels. Hence,

we need a setup that not only requires invariance against affine transformations and articulations but also against partial shapes as well as merged shapes.

Merged or partial shapes pose two problems to shape descriptors and shape matching in general. First, the object boundaries have to be determined and second, the recognition of merged or partial shapes requires shape descriptors or matching methods that can recognize partial models. Existing state-of-the-art methods for petroglyph shape similarity approaches do not fulfill all these requirements. Therefore, we evaluate the discriminative performance of skeletal graphs of petroglyphs. This paper demonstrates, that skeletal graphs are a well-suited description for future part-based retrieval methods.

2 Related Work

Numerous surveys about shape analysis have been published, comprehensive and important in this field are the surveys by Pavlidis [26], Loncaric [21], Zhang and Lu [37] as well as by Yang et al. [35]. Classifications and taxonomies of shape descriptors have been proposed in the mentioned surveys in different variants. A widely used common denominator is the distinction between contour-based and region-based descriptors.

Latecki et al. compared several shape descriptors on the MPEG-7 CE-Shape-1 database [18]. The database contains only complete shapes with closed contours. They propose three main categories for shape descriptors: Contour-based descriptors, region-based descriptors and skeleton-based descriptors. They investigated robustness to scaling and rotation, similarity-based retrieval as well as motion and non-rigid deformations. The weakest performing descriptor in all cases was the skeleton-based approach, the most significant drawback is the lack of robustness against scaling and rotation. The authors assume, that none of the existing approaches to compute skeletons is robust enough. But, since the publishing of this paper in the year 2000, promising skeletonisation algorithms have been proposed and evaluated (e.g. [3]). We summarize region-based, contour-based and skeleton-based approaches that are relevant for our work and include petroglyph-related work where available.

2.1 Region-Based Descriptors

Zhu et al. propose the usage of a slight modification of the generalised Hough transform (GHT) for the mining of large petroglyph datasets [39]. The main arguments for GHT and against other shape similarity measures are the existence of petroglyph images where a single petroglyph consists of several parts and the possibility of merged parts of petroglyphs that drastically change the topology of the petroglyphs. They extensively evaluate their approach and achieve good results. However, they mostly evaluate synthetic petroglyph shapes or simple petroglyph shapes drawn by humans rather than the more exact tracings based on peck marks which we use (see Section 3.1). Deufemia et al. use the radon transform of petroglyph images as shape descriptor for unsupervised recognition

via self-organizing maps (SOM) [9]. In a second step, they use a fuzzy visual language parser to solve ambiguous interpretations by incorporating archaeological knowledge. They evaluate the approach on a large dataset and achieve good results. However, Deufemia et al. as well as Zhu et al. do not consider partial or merged petroglyphs, i.e. part-based retrieval that is necessary for petroglyphs in real-world scenes.

Krish and Snyder propose the shape recognition approach SKS [15] which is based on the generalised Hough transform. They compare the performance of SKS with Hu moments, curvature scale space (CSS, see 2.2) matching and shape context (SC, see 2.2). Besides affine transformations, they evaluate partial shapes. The SKS feature performs good on partial shapes. But, the evaluation data set consists of 31 different shapes only and does not contain merged shapes.¹ Generally, region-based descriptors have the advantage that they do not need complete contours for descriptor extraction.

2.2 Contour-Based Descriptors

Mokhtarian et al. propose curvature scale space (CSS) image matching [25][23][24]. They smooth a contour by convolution with a Gaussian kernel in different scales (i.e. different kernel sizes of the Gaussian kernel). Subsequently, they find the curvature zero crossings on the contour. The descriptor - the CSS image - consists of the zero crossings in a diagram where the size of the Gaussian kernel is on the y-axis and the normalized path length of the curve is on the x-axis. This CSS image is used to match shapes. Mai et al. use the CSS descriptors to acquire contour segments invariant to affine transformations [22]. They utilize the local maxima in the CSS image to locate the affine-invariant points and segment the contour at these points. They match the segments with a dynamic programming approach. They achieve very good experimental results. They outperform the dynamic programming approach by Petrakis et al. [27], who are utilizing contour segments as well. Belongie et al. propose the widely used shape context (SC) descriptor [6]. They sample points on the contour of an object. For each point, they compute the shape context based on the spatial distribution of the other points on the shape contour. They match two shapes by estimating the best transformation from one shape to the other and determining a dissimilarity based on shape context distance, appearance cost and transformation cost.

Deufemia et al. [10] propose a two stage classification of petroglyphs. They use shape context descriptors to provide an initial raw clustering with self-organizing maps. In the second step, they use an image deformation model to classify the petroglyph shapes. They evaluate their approach on a relatively small dataset (17 classes with 3 exemplars) that they enlarge by using 30 affine transformations of each image.

¹ The data set does not include petroglyphs.




Fig. 1. This figure shows a part of a tracing of a rock in Valcamonica (Coren di Redondo, Rock 1). ©Alberto Marretta, used with permission.

2.3 Skeleton-Based Descriptors

Siddiqi and Kimia propose a shock grammar for recognition [31]. Later, Siddiqi et al. propose the usage of shock graphs for shape matching [32]. Shocks are used to provide a structural description of 2D shapes. They are contour-based and deliver a medial axis of the shape, that has additional information for each part of the skeleton. The representation of the shape is a directed acyclic shock graph, which is used for shape matching.

Aslan Skeletons are coarse skeletons [1]. They are matched via tree edit distance and have been evaluated on different data sets with good results [5][12]. However, while they are insensitive to articulation and affine transformations, they can not be used for merged shapes, for shapes with holes and for shapes with large missing parts.

Ling and Jacobs propose the usage of the inner-distance, which is the shortest path between two landmark points (in this case contour points are used) of a shape [20]. Hence, they implicitly embed skeletal information in the descriptor. The distance between two contour points is the shortest path within the silhouette of the shape instead of the Euclidean distance between the two points. They use the idea for three approaches to shape description. First, they combine the inner-distance with multi-dimensional scaling. Second, they utilize the

inner-distance to build a new descriptor based on shape context, and third they extend the second approach with appearance information of the shapes along the inner-distance lines. They evaluate the approach on several datasets with good results. They state, that the proposed descriptors are invariant/insensitive to articulation and are capable to capture part structures.

Bai and Latecki introduce a skeleton-based approach that matches silhouettes based on skeleton paths, which are the geodesic paths between skeleton endpoints [2]. The shortest paths are represented by the radii of the maximum inscribed discs at skeleton points. They use DCE (Discrete Curve Evolution [17]) for the skeleton pruning. The descriptor is on two layers. First, the description of a skeleton endpoint is constructed from the shortest paths starting at this point, and second the similarity of two shapes is computed by matching the descriptors of the skeleton endpoints. They experimentally show, that the method is robust against articulations, stretching and contour deformations. Bai et al. combine contour features with skeletal features to improve shape classification [4]. They state, that contour-based approaches can represent detailed information well, and are up to a certain extent robust against partial and merged shapes but lack invariance against articulation and non-rigid deformation. In contrast to that, skeleton-based approaches are robust against non-rigid deformations. For the contour segments, they follow the ideas of Sun and Super [33], but use DCE to determine the segments. They achieve 96.6% classification rate on the MPEG-7 CE-Shape-1 [18] database. Xu et al. extend the skeleton path approach [2] by also considering junction points for the skeleton paths descriptor [34]. They call the junction points and end points of the skeleton graph critical points. They merge junction nodes based on the paths from the junction nodes to the end nodes. If the sum of path distances of two nodes normalised by the number of end nodes of the graph is below a set threshold the two nodes are merged. They achieve slightly better results than Bai and Latecki [2], and state, that the method is efficient even in the presence of articulation as well as partial and merged shapes. The line of work summarized in this paragraph is mostly built on the skeleton pruning algorithm based on DCE proposed by Bai et al. [3].

To our knowledge, there is no work on petroglyphs that utilizes skeletons or skeletal graphs. We aim at investigating the skeletal graph for petroglyph similarity modeling. In our approach, we use the the promising algorithm by Bai et al. [3] for skeletonisation and skeleton pruning. To investigate the distinctiveness of skeletal graphs, we model petroglyph similarity as graph similarity. In the following, we summarize works that use graph matching to model shape similarity and popular graph matching approaches. There are numerous approaches in shape matching that utilize graphs. Comparable to our approach are methods, that model shape similarity as similarity of the skeletal graph. Klein et al. use tree edit distance to match shapes described by their shock graphs [14]. Di Ruberto uses attributed skeletal graphs to model shape similarity [11]. Aslan skeletal graphs are matched with tree edit distance [5][12]. For our material (see Section 3.1), tree matching is not sufficient, as the skeletal graphs may contain cycles. Hence, we have to use graph matching. There is a long line of work in

graph matching. A recent volume of workshop proceedings edited by Kropatsch et al. includes papers on many aspects of graph-based representations in pattern recognition [16].



Fig. 2. This figure shows examples of the petroglyph dataset that we investigate in this paper. Each column contains examples of one class. We observe, that some of the classes have high intra-class variance of shape. The petroglyphs are from various rocks in Valcamonica. ©CCSP - Centro Camuno di Studi Preistorici and Alberto Marretta, used with permission.

A popular and intuitive similarity measure for a pair of graphs A and B is the graph edit distance (GED). GED defines similarity of two graphs as the minimum number of edit operations (remove node, add node, remove edge, add edge) that are needed to transform graph A to graph B . The computation of the graph edit distance is NP-hard [36]. This problem is addressed in several ways. There are approximations for the GED, e.g. the widely used Hungarian algorithm, or A^* beam search. Another way is to use graph spectra (e.g. [8]) or to embed node and edge attributes of graphs in vector spaces and subsequently use standard similarity measures and machine learning methods to match similarity (e.g. [13] or [19]).

3 Dataset and Approach

3.1 Benchmark Dataset

Manual tracing of petroglyphs on transparent material is the standard documentation technique for pecked rock art. See Figure 1 for a part of a tracing of a rock panel. To obtain such a tracing, the transparent material is placed on the rock panel and each peck mark is traced. We have access to digitized versions of numerous sheets. For this paper, we use a dataset of one hundred tracings of individual petroglyphs in ten classes. Figure 2 shows examples of our material.² Note the high intra-class variability not only in terms of affine transformations and articulations but also in general differences of the shapes in a class. Some classes are perceptually very close, e.g. classes two and three.

² The full dataset is available at <http://ment.org/VISART14>.

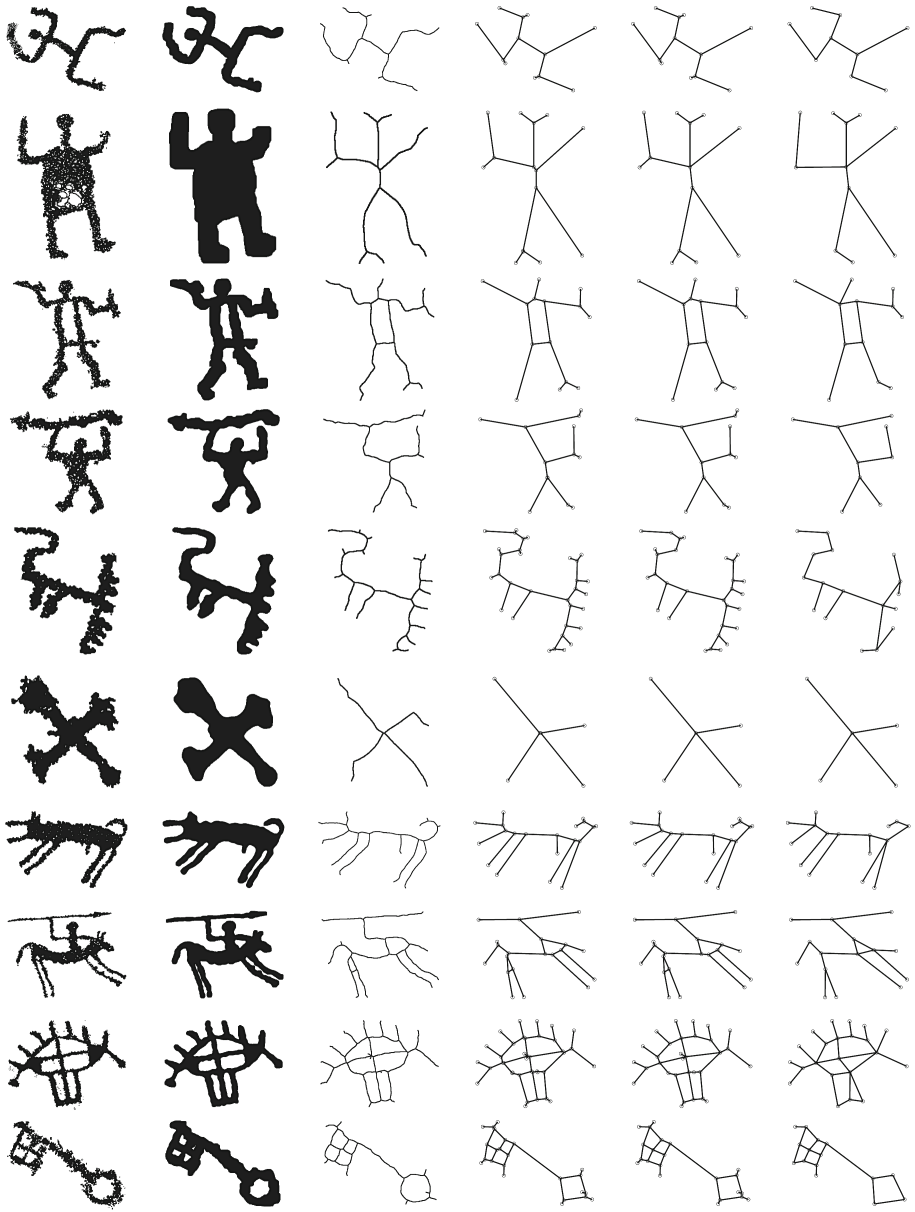


Fig. 3. This figure shows petroglyph tracings (column 1), pre-processed shapes (column 2), extracted skeletons (column 3) and the derived graphs with a pruning threshold of 2, 10 and 30px (columns 4-6). We observe, that the employed skeletonisation algorithm fails to extract details in some cases (heads in rows 1 and 4), while in other cases small details are covered well. Furthermore, we observe, that a high pruning threshold removes skeletal noise (rows 1, 2, 4, 6, 9 and 10), but may also discard relevant topological information (row 5).

3.2 Overview of our Approach

The literature review suggests, that a combination of contour-based and skeletal descriptors should yield optimum results. The existing petroglyph shape retrieval systems utilize region-based features [39][9] as well as contour-based features [10]. These methods are the baseline reference methods we compare our approach with. We concentrate on the investigation of skeletal features for the petroglyph shape recognition problem, because the petroglyph shapes are intuitively already skeletal shapes. Furthermore, we want to use the skeletons as basis for part-based matching in future. We propose to model similarity of petroglyphs as a skeletal graph matching problem. We derive the graphs from the skeletons of the petroglyph shapes. We use each end point and each junction point of a given petroglyph skeleton as nodes for our graph, and the skeleton branches connecting these points as edges. See Figure 3 for exemplary skeletons and derived graphs. We utilize the skeletal graph as descriptor that is invariant to affine transformations as well as articulations. Hence, we discard all spatial information, as we are only interested in topological information. We match the resulting undirected graphs following two strategies. First, we utilize the graph edit distance (GED) as pairwise similarity measure, and second, we use graph embedding (GE) to create a feature vector for each graph and match these feature vectors. We embed the graphs in feature vectors of a normed length by extracting several graph properties (see Table 1), and calculate pairwise distances with Euclidean distance. We classify using a k-NN classifier with an extension for intermediate descriptor fusion.

3.3 Descriptors

We compare the distinctiveness of contour-based as well as region-based shape descriptors with the distinctiveness of undirected skeletal graphs for the petroglyph classification problem. As baseline methods, we utilize a) Shape Context (SC) [6], which has been used by Deufemia et al. for petroglyph classification [10], b) Inner Distance Shape Context (IDSC) [20], as well as c) General Hough Transform (GHT) proposed by Zhu et al. for (GHT) proposed by Zhu et al. for petroglyph classification [39].³ Our proposed petroglyph descriptor makes use of the skeletal graphs in two ways. First, we use the graphs directly as descriptors and measure the similarity with GED. We employ the A* algorithm with beam search and the Hungarian algorithm. Both variants tolerate cycles in the graphs.⁴ Second, we utilize topology features of the undirected graphs for GE.⁵ Please refer to Table 1 for the list of selected features for petroglyph description.

³ For all three descriptors, implementations have been kindly provided by the respective authors.

⁴ We use the implementation in the Graph Matching Framework kindly provided by Kaspar Riesen [28].

⁵ We utilize the MIT strategic engineering tools for network analysis kindly provided by Bounova and de Weck [7]: http://strategic.mit.edu/downloads.php?page=matlab_networks.

We evaluate the distinctiveness of single topology features and of combinations of topology features. In order to be able to use skeletonisation and contour-based descriptors, we need to pre-process our material to achieve continuous boundaries. We use standard filters and morphological operations for this purpose. The resulting shape images are the input for SC, IDSC as well as for the skeletonisation for which we use the method by Bai et al. [3].⁶ We normalize the width of the input petroglyph to 500px and prune the skeletons by joining nodes which have a spatial distance smaller than a threshold. GHT is computed on the original petroglyph images.⁷

Table 1. Topology based features extracted from the skeletal graphs. Each feature is a description of the whole graph. Please refer to [7] for more details.

id	Feature	Description
1	numNodes	Number of nodes
2	numEdges	Number of edges
3	numCycles	Number of independent cycles, also known as the cyclic number of a graph
4	linkDensity	Link density, i.e. ratio of existing links to maximum possible links
5	avgDegree	Average number of links over all nodes
6	numLeafs	Number of leafs, i.e. number of nodes with only one link
7 - 11	histDegrees	5 bin histogram of node degrees, i.e. counts of all nodes with 1,2...5 links. The maximum occurring degree in our dataset is 5
12	sMetric	Sum of degree products across all edges, i.e. for each edge, multiply the degrees of the two nodes connected by the edge and finally sum up the products
13	graphEnergy	Sum of the absolute values of real components of the eigenvalues
14	avgNeighDegree	Average of the average neighboring degrees of all nodes
15	avgCloseness	Average of closeness over all nodes
16	pearson	Pearson coefficient for the degree sequence of all edges of the graph
17	richClub	Rich club metric for all nodes with a degree larger than 1
18	algebConnect	Algebraic connectivity, i.e. the second smallest eigenvalue of the Laplacian
19	diameter	The longest shortest path between any two nodes in the graph
20	avgPathlength	Average shortest path
21	graphRadius	Minimum vertex eccentricity

⁶ We thank Bai et al. for providing the implementation.

⁷ We use resolutions of 10x10px, 20x20px and 30x30px. We report on the best performing resolution 10x10px.

3.4 Descriptor Fusion and Classification

Additionally to the performance of individual descriptors, we evaluate the performance of descriptor combinations. For SC, IDSC, GHT and GED we compute pairwise distances. For the feature vectors we obtain by GE, we calculate pairwise Euclidean distances. The combination of descriptors via the combination (i.e. unweighted summation) of distance matrices would require a normalization in all utilized feature spaces, which could only be provided by heuristically determined thresholds for the maximum dissimilarity that can occur for one specific descriptor. We want to avoid this step, and aim at preserving complementary distinctiveness as far as possible in the classification process. Instead of merging the similarity matrices, we employ single k-NN classifiers ($k=5$) for each individual feature. A straightforward approach would be to combine the classification results of all classifiers by majority voting. To obtain a richer and more expressive basis for making a decision, we refrain from this simple combination of classification results. Instead, we fuse the classifiers at an intermediate step. We use the class labels of the five nearest neighbors of each descriptor and concatenate these to a set of nearest neighbors which contains $5n$ class labels, with n being the number of descriptors combined. Subsequently, we classify according to the conventional rule of k-NN with a majority vote. This incorporates more information than a voting based on the classification results, as we implicitly include the probability of the classification result in form of the number of class members among the nearest neighbors. We validate our results of single descriptors as well as fused descriptors by leave-one-out-cross-validation (LOOCV). We employ accuracy as quality measure, i.e. the ratio of the number of correctly classified petroglyphs to the total number of classified petroglyphs.

4 Experimental Results

4.1 Graph Matching

Table 2 shows the results of GED employed on the unweighted undirected skeletal graphs. The maximum accuracy is 57%. Both employed methods achieve this result at a pruning threshold of 10px. We observe, that the results of both methods tend to decrease with an increasing pruning threshold larger than 10px. We assume, that the distinctiveness of the skeletal graphs first increases, as skeletal noise is pruned, and then decreases after a maximum of 10px as more and more topological distinctiveness is removed in the pruning process (see Figure 3). Furthermore, we observe, that for higher pruning thresholds, A* beam search outperforms the Hungarian algorithm.

Table 3 summarizes the results of the evaluation of single topological features of GE. We observe, that feature 13 (graph energy) yields the best result with 47% accuracy given a pruning threshold of 10px. We achieve also the second and the third best results with a pruning threshold of 10px. This confirms the GED results (see Table 2), where a pruning threshold of 10px delivers the best results as well. Furthermore, we observe, that features 8,10 and 11 (bins 2,4

Table 2. Classification accuracy of GED in percent. We use LOOCV-validated 5-NN classification. The maximum number of open paths for the A* beam approximation is 6000.

	Pruning threshold					
	2	10	15	20	25	30
Hungarian	52	57	57	30	39	37
A* beam	51	57	49	47	47	47

and 5 of the degree histogram) perform around random (10%). We assume that this is due to the fact, that most graph nodes that are not leaves seem to be 3-connected (see Figure 3). Hence, the counts of 2, 4 and 5-connected nodes have weak discriminative capabilities.

Table 3. Classification accuracy (in percent) of GE utilizing single scalar topological features. We use LOOCV-validated 5-NN classification. p denotes the size of the pruning threshold. The three best values are emphasized. Please refer to Table 1 for the descriptions corresponding to the feature ids

p	Feature id																					Avg.
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	
2	25	21	14	35	27	40	40	10	33	10	10	22	25	21	30	23	35	30	27	29	14	25
10	37	44	14	29	31	40	40	9	25	7	10	33	47	23	43	23	29	32	25	31	20	28
15	35	25	17	38	33	37	37	5	21	7	10	30	37	16	30	22	38	30	23	39	23	26
20	31	28	18	31	37	21	21	8	16	10	10	35	35	15	33	26	31	33	29	37	21	25
25	33	32	20	37	39	25	25	12	14	10	10	25	39	13	37	17	37	37	26	33	29	26
30	27	36	13	28	35	23	23	13	8	10	10	33	27	9	36	22	28	34	29	26	28	24
Avg.	31	31	16	33	34	31	31	10	20	9	10	30	35	16	35	22	33	33	27	33	23	

Table 4 reports the GE performance of the best feature combination of each dimension from 1 to 10. We employ brute-force feature selection for each number of feature dimensions, i.e. we evaluate all possible combinations for 1 to 10 out of the 21 single features (see Table 1). We observe, that feature combination strongly improves results. The maximum accuracy is 57%, as it is the case with GED as well.

4.2 Comparison with Baseline Descriptors and Descriptor Fusion

Table 5 contains the results of shape descriptors with which we compare our skeletal graph approach as well as combinations thereof. We observe, that SC and IDSC perform better than GE when employed as single descriptors. GE clearly outperforms the dedicated petroglyph descriptor GHT of [39]. Descriptor fusion generally improves results. The fusion of two or three shape descriptors improves the results slightly. The fusion of all four descriptors improves the results from 82% for the best single descriptor to 88%. This demonstrates, that the descriptors contain complementary information that is well preserved by using the proposed k-NN fusion method.

Table 4. Classification accuracy (in percent) of GE using combinations of 1 to 10 single topology features that perform best for each dimension. We use LOOCV-validated 5-NN classification with brute-force feature selection. The pruning threshold is 10px. The best values are emphasized. Please refer to Table 1 for the descriptions corresponding to the feature ids

Feature ids	Accuracy
13	47
2, 15	54
4, 6, 18	56
6, 7, 18, 20	57
6, 7, 15, 18, 20	57
4, 6, 7, 15, 18, 20	56
1, 2, 4, 6, 7, 15, 17	55
1, 2, 4, 6, 7, 13, 15, 17	55
1, 2, 4, 6, 7, 14, 15, 17, 20	53
1, 2, 4, 6, 7, 13, 14, 15, 17, 20	53

Table 5. Classification accuracy (in percent) of GE, SC, IDSC and GHT and combinations thereof. We use LOOCV-validated 5-NN classification for single shape features and the classifier fusion method described in Section 3.4

Descriptor	Single				Fused													
	x				x	x	x	x	x	x	x	x	x	x	x	x	x	
GE	x				x	x	x	x	x	x								x
IDSC		x			x				x	x		x	x	x				x
SC			x			x			x	x			x	x	x			x
GHT				x			x		x	x			x	x	x			x
Accuracy	57	81	82	39	80	78	54	86	84	83	84	81	86	81	86	81	88	




Fig. 4. This figure shows a sample which is misclassified with GED and GHT. The query image is on the left, and the five nearest neighbors on the right. The first two rows show the result for GED and the utilized graphs. The third row shows the result for GHT

To discuss the limitations of our approach and the weak performance of GHT on our dataset compared to the datasets employed by Zhu et al. [39], we present example query petroglyphs and their nearest neighbors. Figure 4 shows a query petroglyph, that is misclassified by GED as well as by GHT. We observe, that the skeletal graph of the query anthropomorph figure has a topology which is similar to the topology of a cross. Hence, the five nearest neighbors are crosses. The nearest neighbors computed by GHT also fail to determine the correct class for the query petroglyph. We observe, that the spatial distribution of the pixels in the query image is comparable to the nearest neighbors. Figure 5 shows a query petroglyph, that is correctly classified by GE and misclassified by GHT. We observe, that in the case of GE, the four nearest neighbors are topologically very close to the query graph. The fifth neighbor is different. But, we have to take into account, that GE matches with a set of features, that cannot necessarily be understood intuitively (see Table 4). The GHT nearest neighbors show comparable pixel distributions. We assume, that the less competitive performance of GHT on our dataset is related to the fact, that the test datasets used by Zhu et al. are manual transcriptions of petroglyph skeletons (or sometimes outlines, see [39] p95) which leads to simpler shapes than the detailed tracings of peck marks in our material.

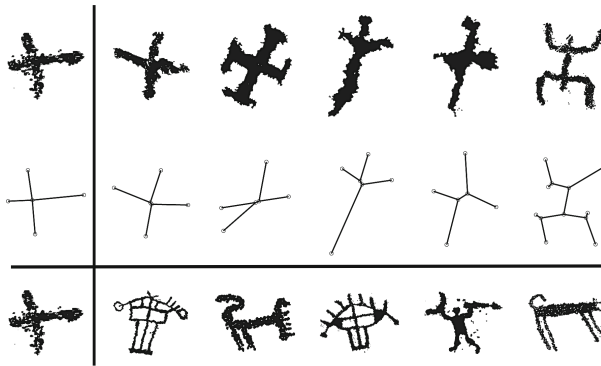


Fig. 5. This figure shows a sample which is correctly classified with GE and misclassified with GHT. The query image is on the left, and the five nearest neighbors on the right. The first two rows show the result for GE and the utilized graphs. The third row shows the result for GHT.

5 Conclusions

We present a novel petroglyph descriptor based on the skeletal graph topology and propose matching with graph edit distance (GED) and graph embedding (GE). For GE, we propose 21 different scalar topological features. We evaluate the descriptor and the matching on a petroglyph dataset containing 10 classes with 10 exemplars and compare the performance with other shape descriptors used in petroglyph classification.

Matching of the skeletal graphs with GE and with GED delivers comparable results. Both matching methods achieve 57% accuracy. GED is of high computational complexity, whereas GE has low computational demand due to low feature vector dimensionality. The two best performing combinations of topology features have only 4 and 5 feature dimensions. The contour-based features SC and IDSC outperform the region-based GHT and the skeletal graph-based GE and GED. GE and GED outperform the region-based GHT. The proposed descriptor fusion clearly improves results. In 5 of 7 descriptor combinations the usage of our descriptor improves results (see Table 5). The combination of our graph-based petroglyph descriptor with other descriptors yields a classification performance of 88% which is not achieved without the proposed skeletal descriptor. This shows that the skeletal features represent information not captured by the contour-based features and the region-based features. We conclude that descriptors derived from skeletons are valuable for petroglyph classification.

Future work will include improvement of our petroglyph descriptor based on the skeletal graph in two ways. First, we aim at improving pre-processing of the shapes as well as skeletonisation. Second, we will investigate the suitability of spatial features. In this paper, we investigated topological information of the skeletal graph, which is an highly invariant abstraction of the skeleton. In future, we will investigate, whether spatial relations of graph parts contain valuable information for our task, because our material has already undergone one step of abstraction by the artists, who created the petroglyphs. Finally, we will make use of our novel petroglyph shape descriptor in the task it was designed for and which it enables: part-based matching.

Acknowledgements. The images of petroglyph tracings used in this paper have been kindly provided by the CCSP - Centro Camuno di Studi Preistorici and by Alberto Marretta, who we thank. The work for this paper has been carried out in the project 3D-Pitoti which is funded from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no 600545; 2013-2016.

References

1. Aslan, C., Erdem, A., Erdem, E., Tari, S.: Disconnected skeleton: shape at its absolute scale. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**(12), 2188–2203 (2008)
2. Bai, X., Latecki, L.: Path similarity skeleton graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**(7), 1282–1292 (2008)
3. Bai, X., Latecki, L., Liu, W.Y.: Skeleton pruning by contour partitioning with discrete curve evolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(3), 449–462 (2007). 00228
4. Bai, X., Liu, W., Tu, Z.: Integrating contour and skeleton for shape classification. In: 2009 IEEE 12th International Conference on Computer Vision Workshops (ICCV Workshops), pp. 360–367. IEEE (2009)

5. Baseski, E., Erdem, A., Tari, S.: Dissimilarity between two skeletal trees in a context. *Pattern Recognition* **42**(3), 370–385 (2009)
6. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**(4), 509–522 (2002)
7. Bounova, G., de Weck, O.: Overview of metrics and their correlation patterns for multiple-metric topology analysis on heterogeneous graph ensembles. *Phys. Rev. E* **85**, 016117 (2012)
8. Demirci, M.F., van Leuken, R.H., Veltkamp, R.C.: Indexing through Laplacian spectra. *Computer Vision and Image Understanding* **110**(3), 312–325 (2008)
9. Deufemia, V., Paolino, L., de Lumley, H.: Petroglyph recognition using self-organizing maps and fuzzy visual language parsing. In: 2012 IEEE 24th International Conference on Tools with Artificial Intelligence (ICTAI), vol. 1, pp. 852–859 (2012)
10. Deufemia, V., Paolino, L.: Combining unsupervised clustering with a non-linear deformation model for efficient petroglyph recognition. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Li, B., Porikli, F., Zordan, V., Klosowski, J., Coquillart, S., Luo, X., Chen, M., Gotz, D. (eds.) *ISVC 2013, Part II. LNCS*, vol. 8034, pp. 128–137. Springer, Heidelberg (2013)
11. Di Ruberto, C.: Recognition of shapes by attributed skeletal graphs. *Pattern Recognition* **37**(1), 21–31 (2004)
12. Erdem, A., Tari, S.: A similarity-based approach for shape classification using Aslan skeletons. *Pattern Recognition Letters* **31**(13), 2024–2032 (2010)
13. Gibert, J., Valveny, E., Bunke, H.: Graph embedding in vector spaces by node attribute statistics. *Pattern Recogn.* **45**(9), 3072–3083 (2012)
14. Klein, P.N., Sebastian, T.B., Kimia, B.B.: Shape matching using edit-distance: an implementation. In: *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2001, Society for Industrial and Applied Mathematics, Philadelphia*, pp. 781–790 (2001)
15. Krish, K., Snyder, W.: A new accumulator-based approach to shape recognition. In: Bebis, G., Boyle, R., Parvin, B., Koracin, D., Remagnino, P., Porikli, F., Peters, J., Klosowski, J., Arns, L., Chun, Y.K., Rhyne, T.-M., Monroe, L. (eds.) *ISVC 2008, Part II. LNCS*, vol. 5359, pp. 157–169. Springer, Heidelberg (2008)
16. Kropatsch, W.G., et al. (eds.): *GbRPR 2013. LNCS*, vol. 7877. Springer, Heidelberg (2013)
17. Latecki, L.J., Lakamper, R.: Shape similarity measure based on correspondence of visual parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(10), 1185–1190 (2000)
18. Latecki, L.J., Lakamper, R., Eckhardt, T.: Shape descriptors for non-rigid shapes with a single closed contour. In: *IEEE Conference on Computer Vision and Pattern Recognition, 2000, Proceedings*, vol. 1, pp. 424–429. IEEE (2000)
19. Li, G., Semerci, M., Yener, B., Zaki, M.J.: Graph classification via topological and label attributes. In: *9th Workshop on Mining and Learning with Graphs (with SIGKDD) (August 2011) (2011)*. 00009
20. Ling, H., Jacobs, D.W.: Shape classification using the inner-distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(2), 286–299 (2007)
21. Loncaric, S.: A survey of shape analysis techniques. *Pattern Recognition* **31**(8), 983–1001 (1998)
22. Mai, F., Chang, C.Q., Hung, Y.S.: Affine-invariant shape matching and recognition under partial occlusion. In: *2010 17th IEEE International Conference on Image Processing (ICIP)*, pp. 4605–4608. IEEE (2010)

23. Mokhtarian, F.: Silhouette-based isolated object recognition through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(5), 539–544 (1995)
24. Mokhtarian, F., Abbasi, S., Kittler, J., et al.: Efficient and robust retrieval by shape content through curvature scale space. *Series on Software Engineering and Knowledge Engineering* **8**, 51–58 (1997)
25. Mokhtarian, F., Mackworth, A.: Scale-based description and recognition of planar curves and two-dimensional shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**(1), 34–43 (1986)
26. Pavlidis, T.: A review of algorithms for shape analysis. *Computer graphics and image processing* **7**(2), 243–258 (1978)
27. Petrakis, E.G.M., Diplaros, A., Milios, E.: Matching and retrieval of distorted and occluded shapes using dynamic programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**(11), 1501–1516 (2002)
28. Riesen, K., Emmenegger, S., Bunke, H.: A novel software toolkit for graph edit distance computation. In: Kropatsch, W.G., Artner, N.M., Haxhimusa, Y., Jiang, X. (eds.) *GbrPR 2013. LNCS*, vol. 7877, pp. 142–151. Springer, Heidelberg (2013)
29. Seidl, M., Breiteneder, C.: Automated petroglyph image segmentation with interactive classifier fusion. In: *Proceedings of the Eighth Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP 2012)* (2012)
30. Seidl, M., Judmaier, P., Baker, F., Chippindale, C., Egger, U., Jax, N., Weis, C., Grubinger, M., Seidl, G.: Multi-touch rocks: playing with tangible virtual heritage in the museum - first user tests. In: *VAST11: The 12th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage - Short Papers*, pp. 73–76 (2011)
31. Siddiqi, K., Kimia, B.: A shock grammar for recognition. In: *1996 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1996, Proceedings CVPR 1996*, pp. 507–513 (1996)
32. Siddiqi, K., Shokoufandeh, A., Dickinson, S., Zucker, S.: Shock graphs and shape matching. *International Journal of Computer Vision* **35**(1), 13–32 (1999)
33. Sun, K.B., Super, B.J.: Classification of contour shapes using class segment sets. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005, CVPR 2005*, vol. 2, pp. 727–733. IEEE (2005)
34. Xu, Y., Wang, B., Liu, W., Bai, X.: Skeleton graph matching based on critical points using path similarity. In: Zha, H., Taniguchi, R., Maybank, S. (eds.) *ACCV 2009, Part III. LNCS*, vol. 5996, pp. 456–465. Springer, Heidelberg (2010)
35. Yang, M., Kpalma, K., Ronsin, J.: A survey of shape feature extraction techniques. *Pattern Recognition* 43–90 (2008). 00163
36. Zeng, Z., Tung, A.K., Wang, J., Feng, J., Zhou, L.: Comparing stars: On approximating graph edit distance. *Proceedings of the VLDB Endowment* **2**(1), 25–36 (2009)
37. Zhang, D., Lu, G.: Review of shape representation and description techniques. *Pattern Recognition* **37**(1), 1–19 (2004)
38. Zhu, Q., Wang, X., Keogh, E., Lee, S.H.: Augmenting the generalized Hough transform to enable the mining of petroglyphs. In: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2009*, pp. 1057–1066. ACM, New York (2009)
39. Zhu, Q., Wang, X., Keogh, E., Lee, S.H.: An efficient and effective similarity measure to enable data mining of petroglyphs. *Data Mining and Knowledge Discovery* **23**(1), 91–127 (2011)