

A study on skeletonization of complex petroglyph shapes

Ewald Wieser¹ · Markus Seidl¹ · Matthias Zeppelzauer¹

Received: 19 August 2015 / Revised: 24 January 2016 / Accepted: 23 February 2016 /

Published online: 2 April 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract In this paper, we present a study on skeletonization of real-world shape data. The data stem from the cultural heritage domain and represent contact tracings of prehistoric petroglyphs. Automated analysis can support the work of archeologists on the investigation and categorization of petroglyphs. One strategy to describe petroglyph shapes is skeleton-based. The skeletonization of petroglyphs is challenging since their shapes are complex, contain numerous holes and are often incomplete or disconnected. Thus they pose an interesting testbed for skeletonization. We present a large real-world dataset consisting of more than 1100 petroglyph shapes. We investigate their properties and requirements for the purpose of skeletonization, and evaluate the applicability of state-of-the-art skeletonization and skeleton pruning algorithms on this type of data. Experiments show that pre-processing of the shapes is crucial to obtain robust skeletons. We propose an adaptive pre-processing method for petroglyph shapes and improve several state-of-the-art skeletonization algorithms to make them suitable for the complex material. Evaluations on our dataset show that 79.8 % of all shapes can be improved by the proposed pre-processing techniques and are thus better suited for subsequent skeletonization. Furthermore we observe that a thinning of the shapes produces robust skeletons for 83.5 % of our shapes and outperforms more sophisticated skeletonization techniques.

Keywords Skeletonization · Petroglyphs · Shape pre-processing · Real-world shape data

✉ Ewald Wieser
ewald.wieser@fhstp.ac.at

Markus Seidl
markus.seidl@fhstp.ac.at

Matthias Zeppelzauer
matthias.zeppelzauer@fhstp.ac.at

¹ Media Computing Research Group, St. Pölten University of Applied Sciences, Matthias Corvinus Straße 15, 3100 St. Pölten, Austria

1 Introduction

In this paper, we present a study on skeletonization of real-world shape data. Skeletonization is a crucial prerequisite for the robust description and indexing of shapes, and for further search and shape retrieval [24]. The investigated data represent manually performed tracings of prehistoric petroglyphs that pose novel challenges to skeletonization due to their complex topology and structure. Petroglyphs are human-made markings on rock surfaces, which were pecked, scratched or carved into rocks [8]. They can be found all over the world and are preserved, studied, and interpreted by archeologists to gain knowledge about early human history. Petroglyphs are an interesting testbed for skeletonization as they exhibit a number of challenges. Depicted motifs range from simple geometric shapes (e.g. crosses) up to compositions of complex hunting, fighting, and dancing scenes.¹ The tracings of the petroglyph shapes may be incomplete due to partial abrasion of the rock surface. Since the petroglyphs are made of individual peck marks they exhibit a complex boundary as well as numerous holes in their interior (see Fig. 1a). Additionally, complex figures may consist of several disconnected parts. Petroglyph shapes can either show filled bodies or just the silhouette of a figure depending on their artistic style. Finally over the years figures have been pecked on top of each other which results in merged shapes.

Petroglyphs are important artifacts that document early human life and development. The digitization and thus permanent preservation of petroglyphs recently gained increasing attention [23, 34]. Recent effort is put into the building of retrieval systems that enable the search for similar shapes as well as the automated classification of petroglyphs into predefined shape classes according to archeological typologies [25]. Following the segmentation of photographs of petroglyphs to get the shapes of the figures [23], the work in this publication is an essential prerequisite for later automated recognition of the shapes based on skeletal descriptors [24, 25].

Existing skeletonization algorithms are not directly applicable to this type of material and yield poor skeletons as shown in Fig. 1b. One reason for the poor performance is that existing methods are usually developed on perfectly segmented shapes with continuous contour and continuously filled regions originating, for example, from public datasets such as MPEG-7 Core Experiment CE-Shape-1 and Kimia-99 [3, 4, 28, 33]. Thus most methods do not fulfill the requirements of noisy real-world data such as that employed in this work. Other algorithms are designed for special tasks (e.g. fingerprint recognition) and therefore rely on specific image properties such as parallel ridges and furrows with well-defined frequency, orientation, and line width [2]. This makes such algorithms inappropriate for our material which show a high variety in composition, line width, and complexity. Aside from different applications (e.g. shape retrieval) enabled by robust skeletonization, the employed petroglyph shapes pose a powerful testbed for the further development of skeletonization algorithms.

Our contribution is a comprehensive investigation of skeletonization in a new application domain that provides noisy real-world data and has rarely been explored so far [29]. We study the applicability of existing skeletonization methods and evaluate their strengths and weaknesses. We identify two major requirements for robust skeletonization: (i) pre-processing of the shapes is necessary to smooth the contour and fill holes in the interior; (ii) existing methods require adaptations to handle the complex structure of the shapes. We propose several improvements of recent skeletonization methods to make them applicable

¹Example: <http://3d-pitoti.eu>.

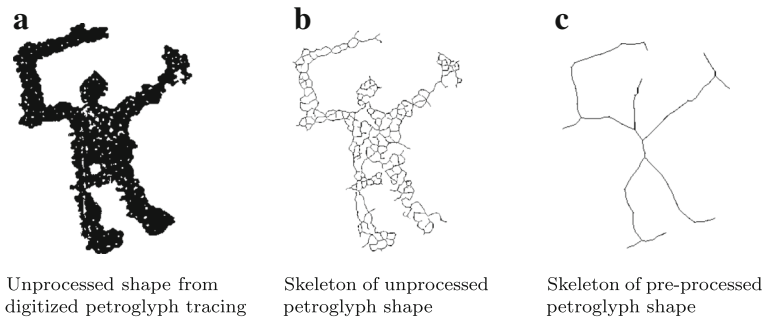


Fig. 1 Comparison of the original petroglyph tracing (a) and its skeleton (b). The petroglyph shape pre-processed with the proposed method results in a significantly simplified skeleton that still preserves the important parts of the original shape (c)

to the investigated material. Finally, we come up with a novel shape pre-processing method that improves the shapes by smoothing and hole filling. The goal of the pre-processing is to reduce the complexity of the shapes and thereby to facilitate the subsequent skeletonization. Our method balances the degree of smoothing and hole filling by an adaptive threshold to yield a suitable tradeoff between preserving the topology and removing noise. An evaluation on our dataset shows that pre-processing is a crucial step for robust skeletonization of noisy and complex shapes like petroglyph tracings (see Fig. 1). Furthermore the proposed improvements of the skeletonization algorithms yield more accurate and complete skeletons.

The paper is structured as follows: In Section 2 we review related work on skeletonization and skeleton pruning and identify suitable algorithms for our task. Section 3 presents our real-world material and its characteristics. We describe our pre-processing approach and the improvements of skeletonization algorithms in Section 4. Experimental setup and results are presented in Sections 5 and 6. Finally we draw conclusions in Section 7.

2 Related work

In this section we review skeletonization and skeleton pruning algorithms, analyze their properties and identify suitable methods for our task. In the literature the usage of terminology for skeletonization is highly ambiguous. Skeletonization, thinning, medial axis transform, distance transform as methodologies and skeleton, medial axis or medial line as their results are used inconsistently [5]. According to Arcelli and Baja [1] algorithms for skeleton computation in *discrete space* can generally be partitioned into two categories: Methods that perform *skeletonization by medial axis transform* produce skeletons following Blum's definition of the medial axis [6] and techniques employing *skeletonization by thinning* derive a thin version of a shape [9, 13]. A third category of approaches performs the medial axis transform to polygonal shapes in *continuous space* [12, 20, 21]. Additionally, there is a group of more recent skeletonization algorithms that utilize physics-based modeling of the shapes [14, 22] for which we suggest a fourth category.

All skeletonization methods are sensitive to boundary noise, i.e. small perturbations of a shape may have large influence on the skeleton (see Fig. 1). To overcome this problem some form of regularization is required [26]. This regularization process is generally referred

to as “*skeleton pruning*”. Shaked and Bruckstein observe that pruning is an essential part of skeletonization algorithms and most recent developments combine skeletonization and skeleton pruning in one algorithm. Skeleton pruning methods can be consolidated in two major categories: The first covers the pruning of skeleton branches based on a significance value calculated for every single skeleton point. This results in a shortening of all skeleton branches. The second class of skeleton pruning algorithms calculate a significance measure for each branch. Based on its significance value a branch is either removed completely or remains in the skeleton [18].

2.1 Point-based pruning approaches

Montanari first develops a form of regularization to detect the most important skeleton branches [19]. He proposes the use of a threshold for Blum’s “*propagation velocity of the wavefront*”. Blum and Nagel extend this idea and propose a boundary/axis weight for the regularization of unwanted branches caused by boundary perturbations [7]. They state, however, that boundary perturbations are not always unwanted distortions but might actually be important features of a shape and therefore pruning should be carried out with great care. Ho and Dyer propose the computation of the relative prominence of a skeleton point by using geometric relations between the maximum generating disk at the point and the contour of the shape [10]. Ogniewicz and Ilg compare several other regularization methods for skeleton points and propose the generation of a skeleton pyramid for further pruning [21]. Telea and van Wijk introduce a skeletonization algorithm based on a fast marching level set method (Augmented Fast Marching Method, AFMM) [31]. For every skeletal point they determine the length of the boundary segment it originates from and prune skeleton points using a single threshold. Howe applies the work of Telea and van Wijk to handwriting recognition using the contour length as saliency measure [11]. Shen et al. compare this and other pixel-based significance measures and introduce a new significance measure for skeleton pruning by calculating the bending potential ratio (BPR) of the contour segment generated by the two points of the maximum inscribed disc that are tangent to the boundary [27]. Telea further improves AFMM by a different saliency metric, and proposes skeletonization for feature preserving shape smoothing [30].

2.2 Branch-based pruning approaches

The methods summarized in Section 2.1 all compute a significance value for each single point of the skeleton. A thresholding of this value leads to a shortening of the branches. Branch-based methods, in contrast, avoid the shortening of branches and instead use a significance value to remove or retain entire branches. Bai et al. [4] propose a novel method for skeleton pruning based on Discrete Curve Evolution (DCE) introduced by Latecki and Lakemper [16]. They determine the contour points of a shape that have maximum curvature and delete all skeleton branches that do not end at one of these points. This approach inspired numerous other state-of-the-art skeleton pruning algorithms. Bai and Latecki further improve DCE by removing the necessity of prior knowledge about the shape [3]. They compute the DCE-skeleton with a fixed parameter (50 vertices) and subsequently add a reconstruction step, which removes skeleton branches with low contribution to the original shape. Yang et al. use the same methodology and extend the reconstruction algorithm to increase speed and to enable the computation of skeletons from shapes with holes [33]. Shen et al. introduce a normalization factor in the reconstruction step that quantifies the tradeoff

between the simplicity of their skeletons and the reconstruction error of the shapes [28]. Liu et al. extract the Generalized Voronoi Skeleton of a shape and then apply DCE to perform a first pruning of the obtained skeleton [17]. Subsequently, they further prune by balancing the visual contribution and the reconstruction contribution of each skeleton branch. Liu et al. further devise a skeleton pruning approach that fuses the information of several different branch significance measures [18]. Recently, Krinidis and Krinidis proposed a new skeletonization approach that smoothes the polygonal approximation of a shape iteratively [15]. In each iteration they determine the most important polygon vertices from the angles of their incoming edges and prune the skeletons by deleting those branches that connect less important nodes.

2.3 Comparison of algorithms

In the following, we specify a number of criteria for the comparison of the skeletonization and pruning techniques presented above.

- **Robustness against remaining insignificant branches:** Insignificant branches are branches that do not contribute essentially to the original shape and thus should be avoided or pruned.
- **Robustness against deletion of significant branches:** A significant branch has an essential contribution to the figure's shape and thus should remain in the skeleton. A deletion would significantly change the structure of the skeleton.
- **Robustness against branch shortening:** Branch shortening occurs when insignificant as well as significant skeleton branches are shortened likewise. This bears the risk of changing the structure of the skeleton.
- **Rotation and scale invariance:** The skeleton of a differently scaled and rotated shape should be equivalent.
- **Number of parameters:** A large number of parameters increases the dependence of an algorithm on user input but at the same time gives more control. We prefer algorithms with a low number of parameters with adequate sensitivity to parameter changes.
- **No prior knowledge about shape needed:** Parameters such as the number of endpoints, or absolute values that depend on the size and complexity of the shape, require a priori knowledge and should be avoided.

Table 1 compares the above presented methods with respect to the specified criteria. The comparison shows that all of the reviewed algorithms differ from each other on at least one criterion. Some of them show major deficiencies such as the deletion of significant skeleton branches, rotation and scale variance, and the need for prior knowledge about the shapes. For our experiments on petroglyph skeletonization we select a subset of the presented algorithms. We perform the selection of methods based on the following considerations: (i) both branch-based and point-based pruning methods should be selected to increase the heterogeneity of evaluated approaches; (ii) all algorithms should be robust to scale and rotation, as the investigated petroglyphs may be rotated and scaled arbitrarily; (iii) the number of input parameters should be low. Based on these considerations, we select the *BPR*-algorithm of Shen et al. [27], the *DCE*-algorithm of Bai et al. [4], and the *SPT*-algorithm of Shen et al. [28]. Furthermore, we add simple morphological *thinning* as an additional method to investigate a skeletonization method that does not apply any pruning.

Table 1 Comparison of recent skeletonization methods with respect to the identified criteria. Note that for point-based pruning approaches the first two criteria do not apply because they do not distinguish between significant and insignificant branches

Algorithm	Robustness against remaining insignificant branches	Robustness against deletion of significant branches	Robustness against branch shortening	Rotation and scale invariance	Number of parameters	No prior knowledge about shape needed
Point-based pruning						
Chord residual + Skeleton pyramid [21]	–	–	no	yes	3	no
AFMM [31]	–	–	no	no	1	no
Boundary length [11]	–	–	no	yes	1	no
Bending potential ratio [27]	–	–	no	yes	1	yes
Saliency metric [30]	–	–	no	yes	1	yes
Branch-based pruning						
DCE-skeleton [4]	no	no	yes	yes	1	no
Discrete skeleton evolution [3]	yes	yes	yes	yes	1	yes
Quick stable skeletons [33]	yes	yes	yes	yes	1	yes
Tradeoff reconstruction error / skeleton simplicity [28]	yes	yes	yes	yes	2	yes
Visual contribution / reconstruction contribution [17]	yes	yes	yes	yes	3	no
Information fusion [18]	yes	yes	yes	yes	1-2	yes/no
Empiric mode decomposition [15]	yes	yes	yes	yes	1	yes

3 Investigated material

The state-of-the-art technique in archeology to document petroglyphs is to manually trace them on transparent foils peck mark by peck mark. Subsequently these foils are scanned and stitched together to large images that cover entire rock surfaces. The tracings show different types of figures and also complex scenes that document ancient life. To extract the individual petroglyph shapes from the digitized tracings, we built a web tool that facilitates the semi-automatic segmentation and manual classification of individual petroglyphs into semantic shape classes according to an archeological typology, e.g. antropomorph, hut, deer, and bird [25]. In a joint effort, a team of archeologists compiled a large data set of more than 1100 petroglyph shapes classified according to two different archeological typologies. Figure 2 shows a small subset of the dataset.

Initial experiments of Takaki et al. showed that skeletonization is a useful abstraction of shapes [29]. Thus it enables higher level applications such as similarity search and automated shape classification, which is our ultimate goal. Petroglyphs pose a challenge for skeletonization as they are made of single peck marks and thus have neither a continuous contour nor continuously filled regions. Figure 2 shows that the shapes have highly varying complexity, contain numerous holes due to incompletely pecked areas, often contain very fine structures (horns of deer, feathers of birds, etc.), and have disconnected parts.

4 Approach

As already discussed in Section 1 the characteristics of petroglyph shapes impede skeletonization which leads to unsatisfactory results (see for example Fig. 1). For robust skeletonization a pre-processing of the shapes is necessary as well as improvements of skeletonization techniques. In the following sections we present a fully automated shape pre-processing method and propose a number of improvements for the selected skeletonization algorithms to make them applicable to petroglyph shapes.




Fig. 2 Some shapes from example classes (e.g. antropomorph, bird, cross, deer, etc.) from our entire dataset of more than 1100 shapes classified according to two different archeological typologies

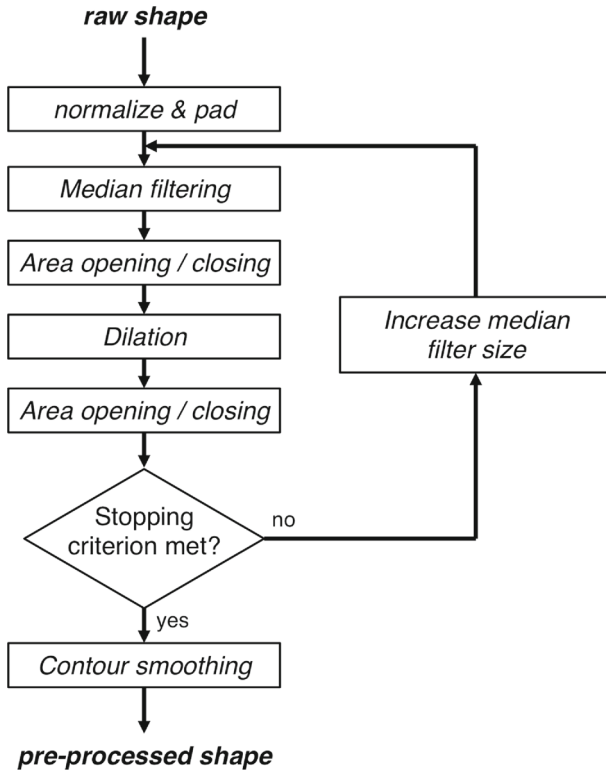


Fig. 3 The workflow of the automated shape pre-processing method

4.1 Adaptive shape pre-processing

The major challenge of pre-processing is to find a solution that forms a good tradeoff across all different types of petroglyphs in our dataset. We define a minimum set of requirements that a suitable pre-processing needs to fulfill to enable more robust skeletonization: (i) close small holes, (ii) smooth the contour, (iii) connect nearby parts, and (iv) avoid the unwanted decomposition of the shapes. We design a method that automatically improves the shapes according to these requirements and that terminates autonomously when a good tradeoff amongst the four aims is achieved. The intuition behind our approach is to combine different structural and morphological filters in a way that they join their strengths. An overview of the method is shown in Fig. 3.

Initially we resize and pad all shapes to normalize the inputs. Next we apply a median filter with size s_{med} to the input. Median filtering removes small holes in foreground and background (salt and pepper noise) and at the same time slightly smooths the contour. Apart from this, however, the median filter may generate artifacts by disconnecting weakly connected blobs. To compensate for these artifacts, we apply an area opening and closing as proposed by Vincent [32]² as well as a dilation operation. We use an area size of t_{aoc} pixels

²Note that the proposed area opening/closing is fundamentally different from a morphological opening/closing as it does not employ a structuring element.

as threshold for area opening and closing, and combine it with a dilation by a disc with a radius r_{dil} to reconnect disjoint parts. We iterate these steps with increasing median filter size until a stopping criterion is met.

The stopping criterion requires a robust indicator function that is suitable for the differently complex shapes in the dataset. We evaluate different indicator functions such as solidity and circularity of the shape, the number and size of foreground and background blobs, and the number of endpoints in the thinning skeleton. Our preliminary experiments show that the most robust criterion is a combination of the number of foreground blobs and the number of background blobs. If the number of both do not change over a certain number of iterations n_{it} , we terminate the pre-processing. The intuition behind this stopping criterion is that if the number of foreground and background blobs remains constant for some time, the figure is likely to be in a robust state where the influence of noise is low.

After the iterative pre-processing has terminated, we smooth the contour by a convolution filter with a Hanning window of size s_{conv} . We smooth the x- and y-coordinates of the shapes' contour points separately which removes contour perturbations and thus reduces the likelihood to get spurious (insignificant) skeleton branches in subsequent skeletonization.

Figure 4 illustrates multiple iterations of the proposed pre-processing method for a given shape. The numbers of foreground and background blobs decrease rapidly during the initial iterations as holes in the shape are removed and closely spaced shape parts are merged. After a few (4) iterations the shape reaches a robust state which is indicated by the fact that the number of foreground and background blobs do not change for a certain number of iterations, i.e. the two indicator functions reach a plateau. We observe further plateaus (e.g. starting at iteration 7, 12, and 15). As iterations proceed the plateaus usually get longer. The stopping criterion finally decides after which plateau length to stop by parameter n_{it} . For the given example a $n_{it} = 3$ would give satisfactory results. After the first plateau the number of foreground blobs increases again and the shape starts to decompose.

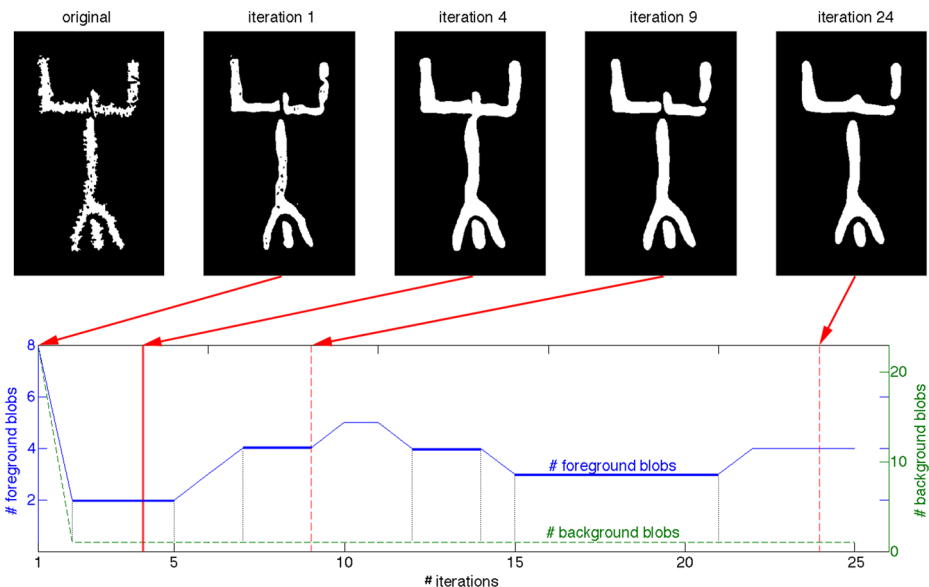


Fig. 4 Sequence of the automated pre-processing of a shape with the proposed method. The numbers of foreground and background blobs decrease rapidly and the stopping criterion aborts the loop at iteration 3. Further iterations would decompose the shape again

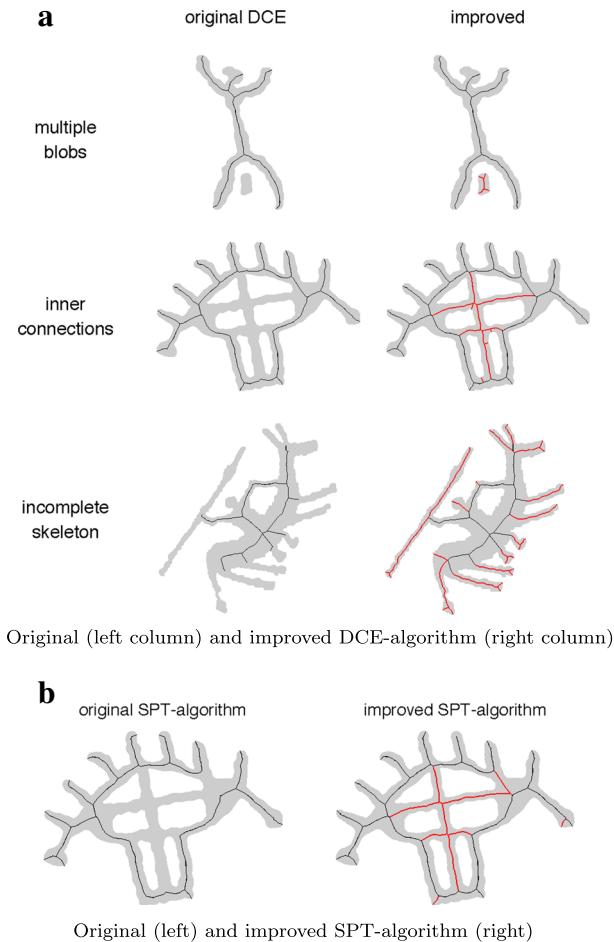


Fig. 5 Comparison of skeletons of original the algorithm (*left* column) with our improved version (*right* column). Our improvements eliminate all shortcomings

4.2 Improvements of existing skeletonization methods

We take the automatically pre-processed shapes as basis for the computation of the skeletons with the chosen algorithms (thinning, BPR, DCE, and SPT). In our preliminary experiments we observe that two of these algorithms (DCE and SPT) suffer from severe problems (see Fig. 5a, left column): The original DCE algorithm [4] is not able to compute the skeleton of multiple blobs in one image, it does not preserve the topology of the shape (loses inner connections), and sometimes generates incomplete skeletons. We propose improvements to compensate for these shortcomings. First, we apply skeleton computation in all blobs of the input instead of on the first blob only (see Fig. 5a, 1st row). To preserve the inner connections, we extend the contour tracing of DCE in a way that all inner contours are traced (see Fig. 5a, 2nd row). The original DCE-algorithm employs either only the concave or only the convex contour points (depending on which set of points is larger) for skeletonization and misses significant shape parts when the concave are selected (see Fig. 5a, 3rd row).

We remove the selection criterion and employ all (convex and concave) points for skeletonization. As a result we obtain complete skeletons even for complex shapes. The effects of our improvements are shown in Fig. 5a, right column. The SPT algorithm [28] utilizes DCE for a coarse pruning of the shape and thus benefits from our improvements as well. A second pruning step in SPT refines the skeleton by deleting branches whose reconstruction contribution to the shape is below a certain threshold. In the original reconstruction step the skeleton path tracing starts at one seed point and traces the skeleton to every endpoint. Thereby inner connections and partly outer connections are missed and thus removed from the skeleton (see Fig. 5b, left). We improve the skeleton path tracing by setting several seed points at the skeleton branch points. As a result all inner connections are preserved in the final skeleton (see Fig. 5b, right).

5 Experimental setup

From the ongoing annotation process of the investigated material described in Section 3, we derive a dataset of 1181 petroglyph shapes to carry out our experiments with.³ In the following, we describe the selection of the parameters for our adaptive shape pre-processing method and the investigated skeletonization methods, and present our evaluation criteria.

5.1 Selection of parameters

For the selection of suitable values for the parameters defined in Section 4.1, we evaluate the pre-processing method on a reduced dataset consisting of 150 representative shapes from the entire dataset. In the absence of a ground truth of pre-processed shapes, we face difficulties in the selection of suitable parameter values. Hence we choose a heuristic approach to estimate robust parameters for the proposed method.

To determine the area opening and closing threshold t_{aoc} [32], we examine the histogram of all blob sizes in the reduced dataset. The histogram is strongly skewed to the left and exhibits a peak between blob sizes of 5 and 20 pixels. Blobs of this size mostly represent noise (e.g. small holes) that usually does not contribute to a figure's shape significantly. With a threshold t_{aoc} between 5 and 20 pixels about 20 % of all blobs are removed and a large portion of noise is filtered out. For the remaining experiments with the entire dataset, we set t_{aoc} to 10 pixels. For the size of the dilation disk r_{dil} we select a rather low value of 3 pixels. This facilitates the reconnection of nearby parts and minimizes the amount of region growing so that the likelihood of merging unrelated shape parts is reduced. From the reduced dataset we observe that typical plateau lengths of the indicator functions are between 2 and 5 iterations. A value of $n_{it} = 3$ iterations yields a good tradeoff between the sufficient smoothing of perturbed shapes, and the decomposition of thin and sparsely connected figures. For some figures a higher value for this stopping criterion leads to a better smoothed shape (see Fig. 6a), whereas for others a higher value leads to a decomposition (see Fig. 6b). To estimate s_{conv} , we apply thinning to the shapes in the reduced dataset and count the number of endpoints. We repeat this process with increasing values of s_{conv} . For filter sizes s_{conv} between 30 and 70 contour points the number of the endpoints remains mostly constant. This shows that the parameter has low sensitivity. We set s_{conv} to 51 pixels for our experiments. The size of the median filter is a non-critical parameter as it is simply

³The dataset can be downloaded at: http://mc.fhstp.ac.at/content/petroskel_dataset.

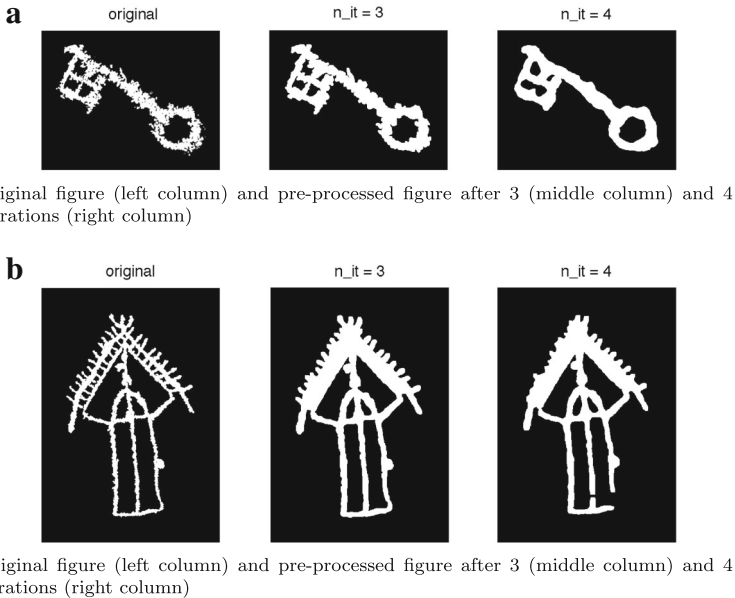


Fig. 6 Evaluation of the stopping criterion for the proposed pre-processing method. For the shape in Fig. 6a a higher value for n_{it} results in a better smoothing, whereas for the shape in Fig. 6b this leads to a decomposition

increased every iteration. We start with a minimum size of $s_{med} = 3$ pixels and increase it by 2 pixels every iteration.

Additional parameters have to be selected for the skeletonization methods. For *DCE* [4], we estimate the parameter for the number of vertices adaptively by counting the number of endpoints of the respective thinning skeleton. For the *BPR*-algorithm [27] and the *SPT*-algorithm [28], we take the parameter values as proposed by their authors. The computation of the *thinning* skeleton is parameter free as it is a simple morphological operation.

5.2 Evaluation

Due to the absence of ground truth shapes and skeletons, we define several perceptual evaluation measures that can easily be judged by a human observer. Subsequently, we evaluate our pre-processing method, and the applied skeletonization algorithms separately on the entire dataset.

For the evaluation of the pre-processing the following criteria are considered:

- All shape parts that are important for visual perception are preserved.
- No independent but closely spaced parts are merged (e.g. legs or feathers).
- Small holes in the shape are closed, disjoint parts are reconnected, and the contour is properly smoothed, i.e. the shape is likely to facilitate subsequent skeletonization.

For the evaluation of the obtained skeletons we apply the following criteria:

- The skeleton preserves the full structure of the shape.
- It exhibits branches for all important parts of the shape.

- It does not have remaining spurious branches.

The evaluation is carried out by a member of the team without any special archeological knowledge. We compile a listing of the full dataset comparing the original shapes, the pre-processed shapes and all skeletons. The evaluator then examines all figures for the criteria defined above and marks problematic figures for further investigation.

6 Results

First, we evaluate the effect of the proposed pre-processing method quantitatively. The intuition behind this evaluation is to obtain an indication of how effectively the pre-processing simplifies the shapes (and thus improves them for subsequent skeletonization). To assess this ability, we utilize a simple observation, namely that typical petroglyph skeletons should usually not exhibit more than 20 endpoints. This can be derived by looking at the shapes in Fig. 2. A consequence of this observation is that skeletons with more than 20 endpoints are likely to contain spurious branches. We use the number of endpoints as a simple heuristic to evaluate the performance of the pre-processing method. For evaluation we count the number of endpoints of the thinning skeletons before and after pre-processing. Figure 7 shows the distribution of the number of endpoints before and after pre-processing. Pre-processing strongly reduces the number of endpoints in the resulting skeletons and maps them to a reasonable range. Thus this evaluation provides a first indication that the proposed pre-processing is beneficial for subsequent skeletonization.

As the employed shape pre-processing operations change the original shapes, for a more comprehensive (qualitative) evaluation of the pre-processing methods a visual inspection of the generated shapes and skeletons is necessary. For this purpose, we manually judge each automatically pre-processed shape according to the criteria presented in Section 5.2. The evaluation results for the proposed pre-processing method on the entire dataset are shown in Table 2.

The pre-processing method performs well for 79.8 % of all figures. Examples are shown in Fig. 8a where unimportant holes in the shapes are removed successfully and the contours are properly smoothed. For the remaining 20.2 % of the shapes we identify different types of errors (see Fig. 8b and Table 2). 9.4 % of all shapes are not smoothed sufficiently and

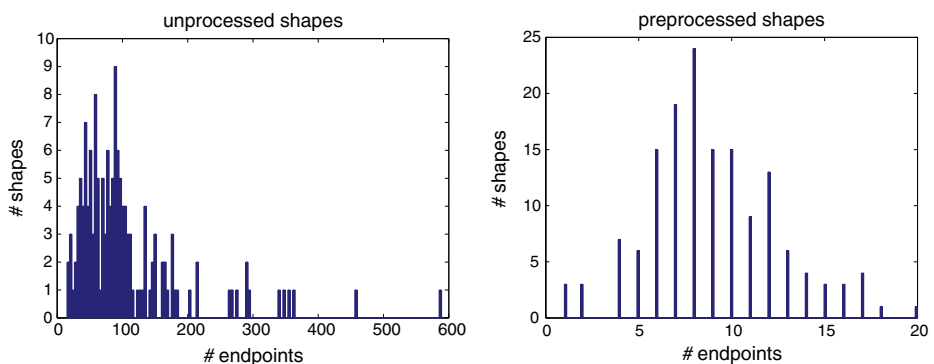


Fig. 7 Comparison of histograms of the number of endpoints of the thinning skeletons of unprocessed shapes (*left*) with the pre-processed shapes (*right*). The number of endpoints significantly decreases and is mapped to a reasonable range

Table 2 Results of the quantitative evaluation of the proposed pre-processing method on the entire dataset

	Not smoothed enough	Parts merged	Details lost	Sum errors	Sum correct
Weak pre-processing	9.4 %	6.9 %	3.9 %	20.2%	79.8%

their contours remain perturbed, which impedes skeletonization. For 6.9 % of the shapes closely spaced parts are merged although it would be important for skeletonization that they remain separated. Only 3.9 % of the shapes are smoothed too strongly and thus important shape parts are lost. The proposed method thus forms a tradeoff between the filtering of noise and removal unimportant shape parts, and the smoothing of important shape parts.

Similarly to the evaluation of the pre-processing, we assessed the quality of the skeletonization methods according to the criteria presented in Section 5.2. The performance evaluation of the skeletonization methods on the pre-processed shapes for the entire dataset is shown in Table 3. On average (over all skeletonization algorithms) we obtained satisfactory results for the majority of our shapes. Figure 9 provides examples of the obtained skeletons and demonstrates that the skeletons are accurate to a high degree. We comprehensively investigated different types of errors that arise during skeletonization and observed that all skeletonization algorithms have deficiencies in certain situations, but none of them completely fails. See Fig. 10 for examples of improper skeletonization.

The skeletons pruned with DCE are in most cases complete (96.3 %), but the algorithm additionally produces a lot of spurious branches (47.7 %). This is related to the fact that the algorithm requires prior knowledge about the shapes (the number of DCE vertices). If this number is set too low, significant branches are deleted even before spurious branches (3.7 %

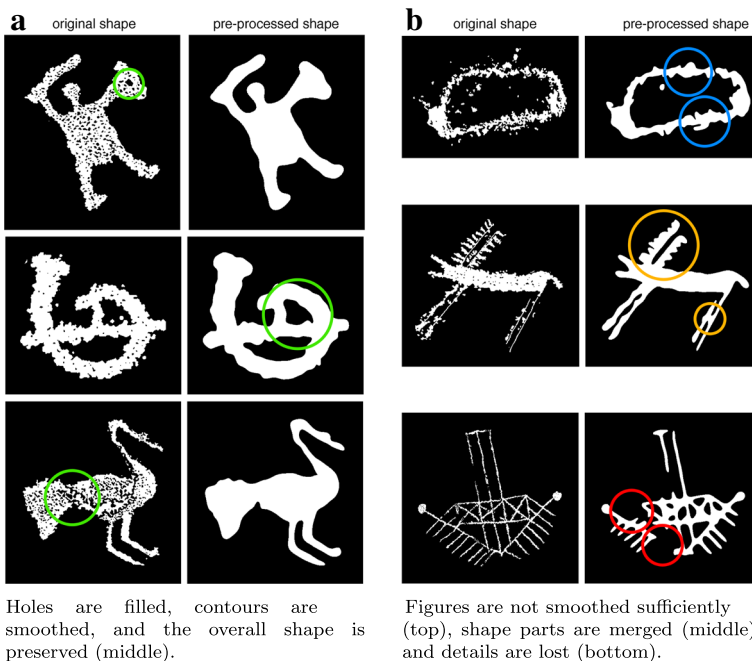
**Fig. 8** Comparison of the results of satisfactory (a) and weak (b) pre-processing of petroglyph shapes

Table 3 Evaluation of the four selected skeletonization algorithms on the entire dataset. BPR and thinning perform best on the pre-processed shapes

Algorithm \ Problem	Spurious branches	Lost parts	Sum errors	Sum correct
Thinning	12.5 %	4.0 %	16.5 %	83.5 %
BPR [27]	6.7 %	6.4 %	13.1 %	86.9 %
DCE [4]	47.7 %	3.7 %	51.4 %	48.6 %
SPT [28]	65.8 %	1.4 %	67.2 %	32.8 %

of all shapes). If the number is set too high, a lot of spurious branches remain (47.7 %). We set this parameter adaptively (depending on the number of endpoints of the corresponding thinning skeleton) since a unique number that is suitable for all shapes does not exist.

SPT builds upon DCE and performs similarly. Although it produces even more spurious branches (65.8 % of all shapes), it deletes spurious branches *before* important ones. Thus for only 1.4 % of all shapes important parts are lost. An advantage of SPT over DCE is that it does not require *a priori* information about the shapes.

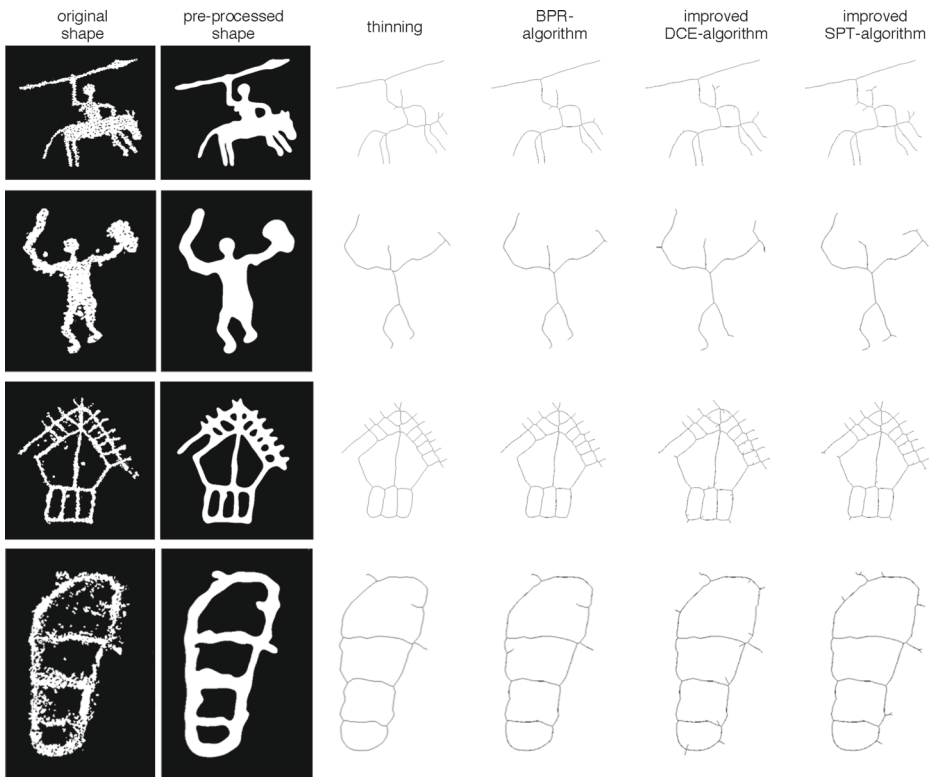


Fig. 9 Examples of successful skeletonization with all selected and improved algorithms based on shape pre-processing with the proposed method. Note that for DCE and SPT a few spurious branches remain in contrast to thinning and BPR

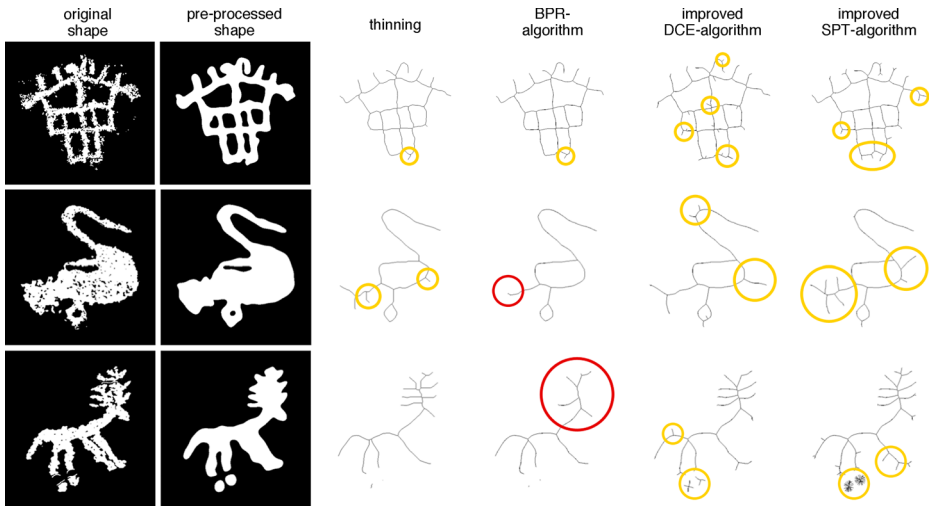


Fig. 10 Examples of partly erroneous skeletons obtained with the selected algorithms. In all depicted examples spurious branches exist (yellow circles) and for some shapes important parts are deleted (2nd and 3rd row, red circles)

The BPR algorithm outperforms DCE and SPT and produces satisfactory skeletons for 86.9 % of all shapes. BPR generates much fewer spurious branches than DCE and SPT (only 6.7 % of all shapes). The BPR pruning is, however, in some situations too strong and thus important branches are removed in 6.4 % of all shapes (see for example tail of the bird in Fig. 10, 2nd row).

A simple thinning results in notably good skeletons for 83.5 % of all shapes which is nearly as good as the performance of the more sophisticated BPR algorithm. Since petroglyph shapes often resemble stick-like figures, they can be well modeled by the thinning algorithm. Additionally, the contour smoothing in the pre-processing avoids the generation of spurious branches by thinning (in 87.5 % of all cases). The results obtained for thinning show that a proper pre-processing can replace an additional skeleton pruning. In only 4.0 % of all cases important branches are missed by thinning.

7 Conclusion

In this paper, we presented a study on skeletonization of petroglyph shapes. We introduced a large heterogeneous dataset of real-world shapes that exhibit numerous challenges to existing skeletonization algorithms and thus poses an interesting testbed. We studied the applicability of existing skeletonization methods and evaluated their strengths and weaknesses. Existing skeletonization methods were developed and evaluated mainly on ideal shapes and are thus not directly applicable to our real-world data. Therefore we improved several skeletonization algorithms to compensate for their shortcomings that became apparent. Additionally, we proposed an adaptive shape pre-processing method that enables the computation of robust skeletons for the complex and diverse shapes under investigation. We performed a large-scale experiment and showed that a proper pre-processing is crucial for the skeletonization of petroglyph shapes. Experiments on skeletonization showed that

pre-processing in combination with a simple thinning yields a good tradeoff for robust skeletonization, whereas more sophisticated skeletonization techniques either generate more spurious branches (DCE, SPT) or delete important ones (BPR). Our experiments clearly demonstrated that the presented pre-processing method and the proposed improvements of recent skeletonization methods solve the additional challenges introduced by our complex and noisy real-world shape data for more than 86% of all investigated shapes.

Acknowledgments Open access funding provided by FH St. Pölten - University of Applied Sciences. 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. This work 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. Further information about the project can be found at <http://3d-pitoti.eu>.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Arcelli C, Baja GSd (1996) Skeletons of planar patterns. In: Kong TY, Rosenfeld A (eds) Topological algorithms for digital image processing, Machine Intell. and Patt. Rec., vol 19. North-Holland, pp 99–143
2. Atul S, Chaudhari ASC (2013) A study and review on fingerprint image enhancement and minutiae extraction. *IOSR J Comput Eng* 9(6):53–56. doi:[10.9790/0661-0965356](https://doi.org/10.9790/0661-0965356)
3. Bai X, Latecki LJ (2007) Discrete skeleton evolution. In: Proceedings of the 6th international conference on energy minimization methods in computer vision and pattern recognition. Springer-Verlag, Berlin, Heidelberg, pp 362–374
4. Bai X, Latecki L, Wy Liu (2007) Skeleton pruning by contour partitioning with discrete curve evolution. *IEEE Trans Pattern Anal Mach Intell* 29(3):449–462
5. Baja GSd (2006) Skeletonization of digital objects. In: Proceedings of the 11th Iberoamerican conference on progress in pattern recognition, image analysis and applications (CIARP06). Springer, DE, pp 1–13
6. Blum H (1967) A transformation for extracting new descriptors of shape. In: Models for the perception of speech and visual form, Proceedings of Meeting held in Boston, Nov. 1964. MIT Press, Cambridge, pp 362–380
7. Blum H, Nagel RN (1978) Shape description using weighted symmetric axis features. *Pattern Recogn* 10(3):167–180
8. Chippindale C, Taçon P (1998) The archaeology of rock-art. New directions in archaeology series. Cambridge University Press
9. Dinneen GP (1955) Programming pattern recognition. In: Proceedings of the March 1-3, 1955, Western Joint Comp. Conf., AFIPS '55 (Western). ACM, NY, pp 94–100
10. Ho S, Dyer C (1984) Medial-axis based shape smoothing. Technical Report 557, University of Wisconsin-Madison Department of Computer Sciences
11. Howe NR (2004) Code implementations by Nicholas R. Howe., <http://www.cs.smith.edu/nhowe/research/code/>
12. Kirkpatrick D (1979) Efficient computation of continuous skeletons. In: Proceedings of the 20th annual symposium on foundations of computer science, 1979. IEEE, San Juan, pp 18–27
13. Kirsch RA, Cahn L, Ray C, Urban GH (1958) Experiments in processing pictorial information with a digital computer. In: Papers and Discussions Presented at the December 9-13, 1957, Eastern Joint Comp. Conf.: Computers with Deadlines to Meet. IRE-ACM-AIEE '57 (Eastern). ACM, New York, pp 221–229
14. Krinidis S, Chatzis V (2009) A skeleton family generator via physics-based deformable models. *IEEE Trans Image Process* 18(1):1–11
15. Krinidis S, Krinidis M (2013) Empirical mode decomposition on skeletonization pruning. *Image Vis Comput* 31(8):533–541
16. Latecki LJ, Lakämper R (1999) Polygon evolution by vertex deletion. In: Proceedings of the 2nd international conference on scale-space theories in computer vision. Springer, pp 398–409

17. Liu H, Wu Z, Hsu DF, Peterson BS, Xu D (2012) On the generation and pruning of skeletons using generalized voronoi diagrams. *Pattern Recogn Lett* 33(16):2113–2119
18. Liu H, Wu ZH, Zhang X, Hsu DF (2013) A skeleton pruning algorithm based on information fusion. *Pattern Recogn Lett* 34(10):1138–1145
19. Montanari U (1968) A method for obtaining skeletons using a quasi-euclidean distance. *J ACM (JACM)* 15(4):600–624
20. Montanari U (1969) Continuous skeletons from digitized images. *J ACM (JACM)* 16(4):534–549
21. Ogniewicz RL, Ilg M (1992) Voronoi skeletons: Theory and applications. In: *Proceedings of the IEEE Conf. on Comp. Vision and Patt. Rec. (CVPR)*, pp 63–69
22. Parker JR (2011) *Algorithms for image processing and computer vision*, 2nd edn. Wiley Publishing, Inc, Indianapolis
23. Seidl M, Breiteneder C (2012) Automated petroglyph image segmentation with interactive classifier fusion. In: *Proceedings of the 8th Indian conference on computer vision, graphics and image processing, ICVGIP '12*. ACM, New York, pp 66:1–66:8
24. Seidl M, Wieser E, Zeppelzauer M, Pinz A, Breiteneder C (2014) Graph-based similarity of petroglyphs. In: *VISART 'Where Computer Vision Meets Art'*, ECCV'2014. Springer, Zürich
25. Seidl M, Wieser E, Alexander C (2015) Automated classification of petroglyphs. *Digital Applications in Archaeology and Cultural Heritage*
26. Shaked D, Bruckstein AM (1998) Pruning medial axes. *Comput Vis Image Underst* 69(2):156–169
27. Shen W, Bai X, Hu R, Wang H, Latecki L (2011) Skeleton growing and pruning with bending potential ratio. *Pattern Recognit* 44(2):196–209
28. Shen W, Bai X, Yang X, Latecki LJ (2013) Skeleton pruning as trade-off between skeleton simplicity and reconstruction error. *Science China Inf Sci* 56(4):1–14
29. Takaki R, Toriwaki J, Mizuno S, Izuhara R (2006) Shape analysis of petroglyphs in central asia. *Forma* 21:243–258
30. Telea A (2012) Feature preserving smoothing of shapes using saliency skeletons. In: *Vis. in Medicine and Life Sciences II*. Springer, pp 153–170
31. Telea A, van Wijk JJ (2002) An augmented fast marching method for computing skeletons and centerlines. In: *Proc. of the symposium on data visualisation 2002*, eurographics ass., VISSYM '02. Aire-la-Ville, Switzerland, pp 251–ff
32. Vincent L (1994) Morphological area openings and closings for grey-scale images. In: *Shape in picture*. Springer, pp 197–208
33. Yang X, Bai X, Yang X, Zeng L (2009) An efficient quick algorithm for computing stable skeletons. In: *Proceedings of the 2nd international congress on image and signal processing (CISP'09)*, pp 1–5
34. Zhu Q, Wang X, Keogh E (2009) 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 '09*. ACM, New York, pp 1057–1066



Ewald Wieser works as junior researcher in the project 3D-PITOTI at the Media Computing Research Group at St. Poelten University of Applied Sciences, Austria. He graduated with Master of Sciences at St. Poelten University of Applied Sciences in Digital Media Technologies in 2014. His research currently focuses on vision computing, machine learning and Deep Neural Networks.



Markus Seidl is working on the technical aspects of interactive systems. His main research interest is Media Processing. Markus has been working for several years in the field of experimental computer art, before he joined UAS St. Pölten as a lecturer. He studied different flavors of Computer Science at the JKU Linz, the University of Vienna, and the Vienna University of Technology, where he received his Master's Degree with distinction. He received the Microsoft Sponsorship award for his master thesis. Since 2011 he is FH-Professor, since 2013 Deputy Head of the IC\M/T - Institute of Creative Media Technologies. Although he is mainly working on the tech side of interactive products, he is a firm believer in the fact, that good interactive products arise from a close collaboration of interface design, user experience and technology. For more information visit his personal homepage at <http://ment.org>.



Matthias Zeppelzauer is a senior researcher at the Institute of Creative Media Technologies (ICMT) at St. Pölten University of Applied Sciences, Austria since 2013. From 2011 to 2015, he worked as a postdoctoral researcher at the Interactive Media Systems Group at the Vienna University of Technology. He graduated with Master of Sciences at Vienna University of Technology in Computer Science and Business Informatics in 2005 and 2006, respectively, and holds a PhD in Computer Science from the Vienna University of Technology. His research focuses on content-based audio retrieval (audio analysis, feature extraction, signal enhancement, noise reduction), visual retrieval (object recognition, model learning), and social media analysis. He is a lecturer for undergraduate- and graduate programs with an emphasis on signal processing and content-based retrieval. Matthias Zeppelzauer received several performance scholarships from the Vienna University of Technology and was awarded by the Austrian Computer Society for outstanding achievements in the area of pattern recognition.