

Direct Elastic Unrollings of Painted Pottery Surfaces from Sparse Image Sets

Peter Houska¹ , Stefan Lengauer¹ , Stephan Karl² , Reinhold Preiner¹ 

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

²University of Graz, Institute of Classics, Austria



Figure 1: Our system takes a set of photographs of antique vases (left) captured from canonical views, as found in many archaeological publications, and computes a complete elastic unrolling of their painted surface area that minimizes projective distortions.

Abstract

An important task in archaeological research is the comparison of painted motifs on ancient vessels and the analysis of their painting style. Ideally, the pottery objects are available as scanned 3D models, from which the painted surface can be unrolled and potential distortions minimized, so that the vase painting and its individual motifs can be directly inspected. Unfortunately, the percentage of digitally captured vessels is small compared to the large body of cataloged photographs. In this paper, we present a method that creates distortion-minimized unrollings of painted pottery surfaces directly from a small set of photographs. We achieve this by exploiting prior knowledge about the data, namely that most objects exhibit rotational symmetry and that strict guidelines were followed when capturing photographs of the ancient vases. Based on the distinctly visible object silhouettes in the photographs we are able to extract proxy geometries of the objects which we encode as per-view geometric maps. By stitching the single-view data, we obtain a combined map capturing the geometry and texture of the entire painted surface. This enables us to minimize typical projective distortions by elastic relaxation. Our pipeline works entirely in 2D image space, circumventing time-consuming 3D scans and surface reconstructions of (often inaccessible) vessels. Using a combination of CPU-based image processing and GPU-based relaxation, results are produced in only a few minutes.

CCS Concepts

• **Computing methodologies** → Image processing;

1. Introduction

Archaeological research in the field of Greek pottery is based on the accurate description and documentation of the vessel shape, but even more importantly, the painted figurative or ornamental decorations on it. Due to its durability, there is an abundance of material for study, which makes this category of physical remains of ancient civilisations exceptionally useful for modern scholars to

address questions of production (dating, workshops and painters, styles), distribution (trade, markets) and consumption (function, use) [Sti02]. Depending on the specific shape, the vase painting is either exhibited only on the front side, both on the front and back side, or continuously around the vase, mostly interrupted by the handles [Ric15]. Using traditional documentation by photogra-



Figure 2: Hand-drawn unwrapping from the 19th century, *Alabastron Brüßel R 224* [LDW58].

phy, it is thus mostly impossible to capture all details of the surface painting in a single, continuous, and distortion-free depiction.

Unwrappings of painted surfaces therefore have a long-standing and well-deserved significance [Wal08]. They show the depiction without dividing it in single images, overcoming the usual photographic distortions as good as possible (Fig. 2). Such unwrappings enable archaeologists to visualize, analyze and interpret the image as a whole as well as in its details. Creating unwrappings manually is time consuming, but more importantly, manual acquisition with tactile tools like tracing paper is often not allowed due to the fragile nature of the surfaces. Peripheral photography introduced in this field in 1965 was soon after abandoned due to its technical complexity and its limitations concerning strongly curved surfaces [Vil65].

With the emergence of 3D scan technologies, algorithms for computer-assisted rollouts were developed to achieve this specific documentation for pottery research. Like photography, 3D scanning has the advantage to be a contactless acquisition method, but has to fulfill the high requirements of texture quality needed in vase painting studies. High-resolution photorealistic models can currently only be practically achieved by using multi-image photogrammetry techniques like structure from motion [Sap18]. Creating useful 3D models needs a lot of (photographic) experience and is time-consuming. More importantly, Greek painted pottery, especially its largest group, the so-called red-figured vases with their inherent homogeneous black glossy background are in most cases incompatible with this technology.

Local efforts in 3D scanning for establishing appropriate documentation for pottery research are mostly neglecting the huge amount of material dispersed over the world and the long research history in this field, whose adequate inclusion is indispensable for comparative studies. A repository of photos of more than 100,000 vases in 26 countries, published in the *Corpus Vasorum Antiquorum* (CVA) since 1922 is fundamental. In addition, more and more image-based databases like the Beazley Archive Pottery Database, the ASCSA Digital Collection or The British Museum collection become openly accessible, enabling researchers to compare large sets of pottery.

Image-based approaches for achieving almost distortion free surfaces are therefore needed in vase painting studies. Fortunately,

a particular characteristic of many available photographs is their structured acquisition procedure, which is determined and documented by the CVA [Blo57, Mom02]. These suggest simple canonical views at orthogonal angles, a near-orthographic projection, and a clear distinction of the vessel from the background. Moreover, most depicted vessels exhibit rotationally symmetric shapes, which in many cases are sufficiently described by their silhouettes alone. These features facilitate a purely image-based geometric reasoning.

In this paper, we propose a novel pipeline that exploits these image characteristics to compute a desired elastic unrolling from a sparse set of input photographs via image stitching, without requiring expensive 3D scanning and reconstruction.

2. Related Work

Previous related work can be found in the area of 3D surface unwrapping and image-based stitching.

2.1. Surface Flattening

Surface flattening or more general unwrapping deals with the problem of creating a flat representation of a 3D surface. Due to the rotational vessel shapes in pottery, computer-assisted rollouts based on 3D models typically utilise proxy geometries like cylinders or cones that exhibit a surface of revolution that best approximates the respective vessel shape [BKMK10, RMK13, MP13]. The mesh is then unwrapped according to the unrolling of these proxies around the rotational axis. Spherical vessel shapes can also be unrolled using a spherical proxy and equirectangular projection [KBMM19]. Today, various tools like the GigaMesh Software framework or CloudCompare allow to perform such rollouts. For general, non-developable surfaces, such flat representations introduce different types of distortions. Especially for pottery shapes with strongly or irregular curved surfaces, the projection can be problematic. Elastic flattening addresses this problem using a numeric optimisation on the vertex locations of an initially flattened mesh [PKBS18]. This method distributes the occurred distortions evenly over the resulting surface, thus achieving results which are very close to hand-drawn unwrappings. To this end, the relaxation solver needs to be provided with the intended distances between the connected vertices of the flattened mesh, which are given by the lengths of edges in the original 3D mesh.

Since our approach does not depend on the availability of a continuous 3D surface, we instead use a purely image-based approach, that extracts the necessary geometric information from the silhouettes visible in each individual photograph. The required distances for the relaxation are encoded in per-image maps, which are stitched along with the image data to obtain a continuous distance profile of the entire surface.

2.2. Image-based Stitching

Stitching is the technique of combining multiple pictures that contain overlapping information. Many image stitching algorithms revolve around the detection of matching pairs of points using Scale-Invariant Feature Transform (SIFT) [Low04], taking care of outliers with the Random sample consensus algorithm [FB81]. Newer



Figure 3: Extracting masks from input photographs.

image-stitching techniques apply deep neural networks to accomplish feature extraction [DTN*20].

After matching pairs of points in single images have been determined, a transformation that aligns these detected features can be defined. In the general case, the set of points in one image cannot be transformed onto the matching set of points on the other without warping one of the images [LWL*18]; otherwise ghosting artifacts, i.e., image regions that cannot be properly matched using a rigid transformation, will occur. Finally, abrupt changes in brightness at stitching seams can be removed [BA83, BL06].

While the above methods all represent relevant tools essential for our work, they are not immediately usable for the automatic computation of complete relaxed surface maps from the sparse image sets available in practice. In the following, we describe our new image-based approach that computes such elastic unrollings automatically from a few canonical images.

3. Our Method

Our method comprises five major stages. First, masks are extracted from the input photographs, and handles are removed from them to provide rotation symmetric shapes (Section 3.1). Silhouettes found in the masks are then used to perform a cylindrical unwrapping of each input photograph, as well as a computation of their directional distance profile across the 3D surface (Section 3.2). The unwrapped images and distance maps are then stitched to a complete flat texture mosaic (Section 3.3). In this mosaic, we identify the relevant motif region and crop it to extract an image of the corresponding area of interest (Section 3.4). The resulting image is relaxed to minimize remaining distortions (Section 3.5).

3.1. Mask Extraction and Handle Removal

The first step in our pipeline is the extraction of the object from the background in each supplied photograph, which is then encoded in a binary mask [LKL*20] (Fig. 3). The silhouettes of these binary masks are used for estimating the object's 3D shape and with it the projective deformation of the surface motifs. This is possible because the vessels resemble bodies of revolution and their silhouette contours provide sufficient information for an approximative 3D shape estimation (Section 3.2). The principal feature contradicting this assumption is the potential presence of handles on the vessels, as shown in Fig. 5. In this step of our pipeline, we therefore remove all handles from the image masks in an automatic processing




Figure 4: Handle Removal. (a) Left and right gradient curves dc_l , dc_r of the milled silhouettes, normalized by image width. (b) The low-pass filtered difference of the gradient curves is used to find suitable bounds a and b for silhouette reflection.

stage. Banterle et al. [BID*17] remove handles from a single 2D vase contour, drawn with clear inner and outer silhouette curves. In contrast, we need to remove handles from binary masks providing a single contour, while maintaining the approximate symmetry across a possibly asymmetric shape. To this end, our approach is to reconstruct the silhouette on the side of the handle from the other side for one-handed masks, while double-handled vessels are treated via feature reduction. Our method works in three stages:

Silhouette Milling. In the first stage, we extract the list of foreground intervals for each row of the binary mask. In rows with multiple intervals, as given in the presence of handles, we only keep the interval that contains, or is closest to, the horizontal image center, and overwrite the remaining intervals by background pixels (red regions in Fig. 4a). At the same time, the number of intervals to the left and to the right of the remaining central intervals are accumulated over all rows in two counters n_l and n_r . We then define the handedness of the vessel based on the ratio $H = n_r / (n_l + n_r)$, as

$$handedness = \begin{cases} \text{left} & \text{if } H < 0.4 \\ \text{right} & \text{if } H > 0.6 \\ \text{double} & \text{otherwise.} \end{cases} \quad (1)$$

Finally, we compute the gradient curves dc_l and dc_r of the remaining foreground interval boundaries along the vertical axis via pixel-wise differentiation (orange curves in Fig. 4a).

Handle region computation. The remaining task is to remove the acute residual features remaining after milling, and restore a symmetric mask. For one-handed vessels (handedness *left* or *right*), we therefore first determine the interval $[a, b]$ along the vertical axis that deviates from the handleless side. To this end, we analyse the (low-pass filtered) difference $\delta = f_{lowp}(dc_l - dc_r)$ of the gradient curves (orange curve in Fig. 4b). In the first step, δ is cleaned by clamping all values with $|\delta_y| < \delta_{min}$ to zero (dark blue curve in Fig. 4b). Intermediate bounds are then found by scanning the clamped function $\tilde{\delta}$ from the top and the bottom until the first




Figure 5: Variety of vessel masks with highlighted handles removed by our approach.

non-zero value is hit (dashed bounds in Fig. 4b). The intermediate bounds are then padded by scanning downhill along the original curve δ until a local minimum is found (thick red bounds). We used a threshold of $\delta_{min} = 0.1\%$ of the image width in our example.

Silhouette reflection. In the final step, the boundary curve of the handleless side is reflected over to the side of the handle within the determined vertical interval $[a, b]$, depending on the handedness of the object. For left-handed masks, c_l is updated by

$$c_l(y) = (c_r(y) - a_r) \cdot \frac{b_l - a_l}{b_r - a_r} + a_l, \quad y \in [a, b] \quad (2)$$

where a_l, b_l and a_r, b_r are the interval boundary values of c_l and c_r , respectively (Fig. 4b). The procedure is symmetric for the right-handed case. Figure 5 depicts a diverse set of handle shapes that are removed by our technique.

3.2. Image-based Unwrapping

Inspired by Karras et al. [KPP96], we generate a cylindrical unwrapping based on a virtual proxy geometry obtained from the object's silhouette for each object depiction I_{orig} . To this end, we leverage the fact that object depictions in archaeological publications are captured in a standardized manner [Mom02, Blo57, Sch14], showing the objects from canonical angles with a near orthographic projection. Moreover, we assume objects to be rotationally symmetric (apart from handles), which is an intrinsic feature to most ancient pottery artifacts due to their manufacturing on a potter's wheel.

As a first step, object silhouettes are extracted by applying a contour detection algorithm [S*85] on the respective mask image. Subjecting the derived contour to a PCA yields the location and direction of the object's rotation axis, which is assumed to run through the mean point and in direction of the first principal component. Secondly, we rotate the object image and its mask such that the rotation axis aligns with the Y-axis of the image coordinate axes. We use the rotation axis to split the contour into a left and a right part




Figure 6: (a) Object depiction with extracted silhouette, (b) virtual proxy geometry, (c) unwrapped object surface.

(Fig. 6a). Let $L = \{l_y\}$ and $R = \{r_y\}$ be the absolute distances of the left and right part respectively, for offsets $y \in [y_{min}, y_{max}]$. We define the geometry of a virtual proxy model such that parallel circular segments span from L to R around the rotation axis. At offset y , the radius of the arc

$$r_y(\varphi) = \frac{r_y - l_y}{\pi} \varphi + l_y \quad (3)$$

is defined such that it linearly scales from l_y to r_y over the azimuth angle $\varphi \in [0, \pi]$. As target image of the projection we define a pixel grid of dimensions $w \times h$, with $w = \pi \cdot \max L \cup R$ as the half-perimeter of the wrapping cylinder of the object's solid of revolution, and $h = y_{max} - y_{min}$. Each pixel $p = (p_x, p_y) \in \mathbb{R}^2$ of this image corresponds to a position on the surface of the virtual 3D proxy $q(p_x, p_y) = (\frac{w}{s_y} \cos \varphi_y, p_y, \frac{w}{s_y} \sin \varphi_y) \in \mathbb{R}^3$ (Fig. 6b). The azimuth angle $\varphi_y(\frac{s_y}{w} p_x)$ is given by

$$\varphi_y(x) = \begin{cases} \frac{x}{l_y} & \text{if } l_y = r_y \\ -\frac{l_y}{\kappa} \pm \sqrt{\left(\frac{l_y}{\kappa}\right)^2 + \frac{2x}{\kappa}} & \text{otherwise,} \end{cases} \quad (4)$$

with $\kappa = \frac{r_y - l_y}{\pi}$, $s_y = \frac{\pi}{2}(l_y + r_y)$ is the arc length of the arc at offset y and is obtained by solving the integral over Eqn. (3) from 0 to π .

From the map of corresponding 3D positions $Q = (q) \in \mathbb{R}^{w \times h \times 3}$ three outputs are computed. First, an RGB image I_{flat} of the unwrapped surface (Fig. 6c), which we obtain from a pixel-wise linear interpolation. The RGB values of $I_{flat}(p)$ are given by

$$I_{flat}(p) = \frac{I_{orig}(\lfloor q_x(p) \rfloor, q_y(p))}{|\lfloor q_x(p) \rfloor - q_x(p)|} + \frac{I_{orig}(\lceil q_x(p) \rceil, q_y(p))}{|\lceil q_x(p) \rceil - q_x(p)|}. \quad (5)$$

Secondly, the Euclidean distances $N = (n_p[i]) \in \mathbb{R}^{w \times h \times 8}$ of 3D positions Q between 8-neighboring pixels are computed and stored in 8 individual distance maps per image (Fig. 7).

3.3. Stitching

The next stage in our pipeline involves stitching the individual cylindrically unrolled photos to a single complete depiction. We




Figure 7: Four out of 8 distance maps of a given image I_{flat} , corresponding to different neighbor directions. Each distance map i stores the Euclidean distance $n_p[i] = \|q(p) - q(p_i)\|$ between corresponding surface points of each pixel p and its i -th neighbor p_i .

employ the stitching algorithm by Li et al. [LWL*18], which is based on robust elastic warping. The method searches for corresponding feature points between two neighboring images, and then applies a Bayesian model to remove local outliers, which makes feature detection, and thus the entire stitching, more robust. Once all corresponding feature points are found, the images to be stitched are turned into regular grids of vertices and deformed to match up the feature points, which completes the stitching.

Even though the given implementation is capable of automatically determining the horizontal arrangement of multiple images to be stitched, we found that the stitched mosaic is prone to vertical drift. We therefore decided to perform a hierarchic pairwise stitching, providing a manually determined image sequence. To this end we compute a similarity distance of all pairwise combinations of input images based on SIFT key points and the ratio test [Low04]. Starting from the pair with the lowest distance we greedily append all remaining images to the left and right, yielding the resulting image sequence.

Given the correct order of images, we can provide the stitching algorithm with prior knowledge of the central image, which allows to reduce possible vertical drift from successively warping stitched image pieces. Additionally, we redundantly append the rightmost picture also on the left side of the mosaic, which allows for an extended surface depiction around the seam of the cylindrical unrolling and thereby reduces the risk of image boundaries cutting through important motifs on the painted surface. To avoid smaller objects being wrongly interpreted as being present twice on the original surface, redundant areas are later visually highlighted in the final surface presentation. This is realized by mapping duplicated surface areas to grayscale colors (Fig. 1b). When working with black-and-white photographs, these redundant areas could also be changed in brightness, or tinted with color instead.

Besides the cylindrically unrolled photographs, we also need to stitch their associated distance maps to yield a complete continuous map of directional surface distances required by the subsequent relaxation. To this end, the exact same transformations that align the individual photographs to a large surface mosaic are applied to the individual per-image distance maps (Fig. 7).




Figure 8: (a) Stitching of individual cylindrical unwrappings. Color boundaries indicate different views. (b) Multiresolution spline blending of the registered unwrappings [BA83].

Figure 8a shows the result of stitching four cylindrically unwrapped input images, where the rightmost image is appended a second time on the left. Note how warping is applied to each individual image to match up all images.

In a final step, abrupt changes in brightness at stitching seams are removed by applying Laplacian multi-resolution blending [BA83] to the warped image mosaic. Figure 8b shows the final, clean result of the stitching phase.

3.4. Motif Cropping

After a complete mosaic of the unwrapped surface is obtained, we apply automatic vertical cropping to the relevant area of the image that contains the motifs of interest. Many Attic vases depict motifs only within a narrow vertical area, while large areas around the neck, shoulder or the lower part of the vessel are covered with monotone black or red bands. In manual drawings or off-tracings for domain publication and documentation, archaeologists only focus on the motif regions and bounding ornaments, and virtually crop the remaining monotone areas as well. Moreover, vertically cropping the area of interest can have a significant influence on the result of the subsequent elastic relaxation step (Section 3.5), as the removal of bulky or tapered surface areas in the mapping will strongly reduce the induced surface strain and with it the resulting residual deformation of the final depiction.

As the motifs on Attic black- or red-figured vases are dichromatic by nature, we propose to identify the vertical region of inter-




Figure 9: Motif cropping. (a) Binarized version of Fig. 8b (horizontally squeezed for illustration). (b) Monotony curve along the y-axis and automatically determined area-of-interest boundaries (dashed black). Red lines in (a) define the padded cutting region.

est by analyzing the vertical contrast profile of their texture. To this end, we first binarize the horizontally unwrapped surface mosaic using Otsu’s thresholding [Ots79] (Fig. 9a). We then compute an intermediate vertical monotony curve

$$m(y) = |\#b_y - \#w_y| / \text{width} \quad (6)$$

that encodes a monotony $m \in [0, 1]$ based on the number of black and white pixels $\#b_y$ and $\#w_y$ in each row y . Being defined on a per pixel row, this curve can exhibit very high frequencies, e.g., in the presence of thin gaps in the motif arrangement (orange curve in Fig. 9b). We therefore use a median filter of certain kernel size (3% of the image height in all our examples), yielding a final filtered monotony profile $\hat{m}$. Finally, starting from the vertical center of $\hat{m}$, the bounds of the contrast-rich motif area are found by scanning upwards and downwards until a monotony threshold $m_{\text{bound}} = 0.9$ is reached. Fig. 9b illustrates $\hat{m}$ as blue curve, where values $\hat{m}(y) > m_{\text{bound}}$ are clamped to 1 for better depiction of the monotony plateaus bounding the motif area. The texture mosaic and its corresponding stitched distance maps are then cropped after adding a user-defined padding distance to the resulting bounds (red lines in Fig. 9a). As the stitching can potentially produce a texture mosaic with irregular boundaries, the image is further cropped both horizontally and vertically to remove any residual background that is not part of the texture mosaic.

3.5. Elastic Relaxation

The last stage of our pipeline takes the cropped texture mosaic from the previous stage and applies an elastic relaxation [PKBS18] in order to minimize length distortions over the entire image. To this end, the image is resampled into a textured triangle mesh, where each quad of two triangles covers an area of 16×16 pixels. The average mosaic resolution produced in our experiments was about 2000×1000 pixels, resulting in a mesh of around 8000 vertices. Upon this mesh a mass-spring system is defined, whose springs connect the 8-neighborhood of vertices.




Figure 10: Final complete distance maps for all 8 directions, filtered using a 7x7 Gaussian kernel.

To allow the relaxation to optimize for an elastic deformation of minimal distortion, we have to define the intended rest lengths r_{ij} for each spring connecting vertices v_i and v_j . In a classic relaxation process of a flattened 3D mesh, these rest lengths are defined by the corresponding edge lengths in the original 3D mesh [PKBS18]. However, as we have circumvented the expensive reconstruction of a continuous 3D surface, we instead obtain the necessary distance information from the distance maps stitched and cropped along with the surface texture throughout our pipeline (Figure 10). To determine the rest length r_{ij} for the directional spring of a given direction $i = \{0, 1, \dots, 7\}$, we integrate the directed distance values stored along a raster line in the i -th distance map between the pixel locations corresponding to vertices v_i and v_j .

Applying the relaxation solver on the texture mosaics of Greek vases typically results in an upward bending of the surface due to a contraction of the upper shoulder and neck regions of the vessel surface (right column in Fig. 11).

4. Results and Discussion

In the following, we evaluate our method using a variety of photos from archaeological vessel publications, exhibiting different shape, motif and quality characteristics. All results have been produced on a consumer-level laptop with a 2.4 GHz Core-i5 6200U CPU, 16 GB RAM, and HD Graphics 520 integrated graphics processor.

4.1. Implementation and Evaluation Data

All stages of our pipeline are realized as CPU scripts, except the relaxation stage, which uses a GPU-based solver [PKBS18]. For the stitching stage we have employed the publicly available Mat-Lab implementation of elastic warping [LWL*18]. Data transfer between pipeline stages is realized through intermediate files that are stored on disk. Using this setup, our entire pipeline takes about 4 to 5 minutes to produce a final result from four 500×1000 input photos.

4.2. Unrolling Sets of Canonical Photographs

The primary use case we target is the computation of a complete elastic unrolling from a small set of photographs that have been



Figure 11: Unrolling results for sets of multiple photographs. Three different kinds of the shape lekythos.

taken according to standard CVA acquisition guidelines. To document circumferential paintings on vessels, the CVA typically exhibits four photographs captured from canonical views. These typically exhibit best possible lighting conditions, orthogonal viewing angles and a near-orthographic projection due to large focal lengths. Figure 1 shows such a set of images, and the unrolling resulting from our approach. Several further results are presented in Figure 11. Here, the left column exhibits vessels with a single handle on the top, black-figured motifs of different complexity, and differently-sized monotone regions which domain experts intend to remove from the depiction. The center column shows the corresponding result of the stitched cylindrical unwrappings (Section 3.3). The yellow bars indicate the area of interest that is automatically determined by the motif cropping stage (Section 3.4). The right column shows the intended result after relaxing the cropped mosaic. Redundant surface areas at the boundaries of wrapped-around motifs are highlighted by desaturation. For the lekythos in Figure 11c only three canonical photographs have been available, and therefore no redundant surface area is visible. In this example,

unpainted background areas could also be cropped horizontally for a better distortion minimization of the actual motif. However, the absence or extent of motifs can also be a valuable information for vase painting studies, and should therefore be captured in the complete unrolling.

In general, we can observe that the elastic unrollings exhibit an overall reduced distortion. This is particularly visible at motifs on the more narrow shoulder regions, such as the dragon (Fig. 11a), the floral ornaments (Fig. 11b), and the athletes (Fig. 11c), which all appear (partially significantly) stretched in the stitched cylindrical unwrappings shown in the center column.

4.3. Single-Picture Unrollings

Apart from the stitching of multiple images, strong projective distortions can already appear in single photographs, which are otherwise sufficient to cover the extent of a motif. An example is given in the amphora in Figure 12a, depicting a youth whose head is almost imperceptible due to strong perspective shortening in the



Figure 12: Results from single photographs, where the stitching phase has been omitted.

shoulder region. The automatic elastic unrolling is able to largely recover its original shape and even depict nuanced features like the ear or eye. Distortions are also reduced along the horizontal axis as shown in Figure 12b. Here the cylindrical unwrapping already allows a clearer view on the fillet held by the Nike. The motif cropping identifies the upper boundary above the shoulder edge due to residual patterns in this area. In contrast, experts would typically cut the motif at the geometric discontinuity in such a case. However, the relaxation result preserves the shoulder decoration as well as the presence of the edge, which is visible in the strong stretch in the upper area and characteristic for a lekythos. Moreover, it gives a cue about the connection between the decorations on the vessel's shoulder and its body.

4.4. Influence of Viewing Angle

In the presence of strong perspective shortenings, the CVA acquisition standard proposes to deviate from the optimal viewing angle to allow a better view on its motifs. Moreover, many photographs of such vessels, e.g., available on public picture repositories, do not always provide ideal orthogonal views. These aspects raise the question how view angle deviations influence the performance of our method. To this end, we compare two photographs of the same vessel captured from both an ideal orthogonal angle and a slanted angle (Fig. 13). The motif on this hydria depicts two figures, Triptolemos and Kore, whose heads are again almost imperceptible under an orthogonal view. Compared to the youth in Fig. 12a, their heads and upper bodies exhibit a more prominent bending towards the image center, resulting in a more complex distortion. Given the orthogonal view photograph (Fig. 13a), our relaxation result restores their




Figure 13: Graz G 30 (Triptolemos and Kore)

faces to a certain extent, but some residual distortions and blurring is noticeable. In contrast, the slanted view in Figure 13b naturally better captures the heads of these figures. Now the silhouette exhibited in the photograph does not accurately reflect the true shape of the vessel anymore, which leads to expectable distortions in the computed distance maps. These are visible in clear residual distortions of Triptolemos' sceptre, which our method managed to correctly straighten up in the orthogonal view. However, the overall image remains intact, and in particular, provides a clearer depiction of the area around their heads. This shows that our system tolerates moderate deviations from the perfect orthogonal view, which allows for improved flexibility in practical applications.

4.5. Limitations

Our method is based on two major assumptions. First, we assume that vessels are rotation symmetric around their upward axis. While this assumption holds for the vast majority of the cataloged corpus of ancient pottery, a small fraction of objects exhibits asymmetric silhouettes, either due to their inherent shape or due to fracturing. Secondly, we require the input pictures to be captured at sufficiently large focal length to provide sufficient overlap for the stitching in the presence of strong curvatures.

Figure 14a shows a set of photographs of an alabastron that violates both these assumptions. First, the photographs exhibit a strong bulky shape painted down to their bottom, which is not visible in its entirety as the pictures were captured from a close distance. Secondly, the vessel has a broken mouth that results in an asymmetric silhouette shape at its upper end. Since the vessel exhibits contrast-rich surface paintings from its base up to its mouth, neither of the two regions are cropped and are therefore subject to both the stitching and the relaxation. At the top, the asymmetric region around the



Figure 14: (a) Problematic case considering bad geometry (asymmetrically broken mouth) and non-conform acquisition. (b) Ideal orthographic views of a digital 3D model with asymmetric regions cropped, resulting in an ideal unrolling (c). (d) Reference result based on the cylindrical unwrapping of the entire 3D model [PKBS18].

mouth leads irregularities in the distance map, resulting in a non-symmetric rollout at the upper region. At the bottom, insufficient overlap leads to a stitching error, visible in a discontinuity at the tail of the Typhon. For comparison, we have created four *virtual* photographs of a 3D reconstruction of this object, using an ideal orthographic projection and with the asymmetric mouth removed (Fig. 14b). Given these input images, our method produces a symmetric and artifact-free result (Fig. 14c). Figure 14d shows a classic mesh-based elastic flattening of the same 3D model [PKBS18].

Another aspect that affects the quality of the input data is the condition and clarity of the surface motifs. The example in Figure 9a shows that the motif cropping does not only detect motif areas, but can also wrongly classify noise, disturbing highlights, or secondary image features like handles as motifs. Hence, cropping can be overly conservative or even fail to crop anything at all. Conversely, bad lighting conditions or low-contrast images can lead to a corrupt binary image, possibly cropping the entire image and passing on invalid data to the relaxation stage.

5. Conclusion and Future Work

We have presented a novel, purely image-based algorithm for generating distortion-minimized surface unrollings from just a few photographs, which are taken from archives of archaeological institutions that are readily available in select CVA volumes [ZE18]. The core insight facilitating our method is that the rest lengths required for the relaxation phase do not need to be extracted from a full 3D reconstructed surface of the vessel, but can be directly extracted from an implicit shape estimate of each input image and stitched together along with the surface textures. Our method considerably lowers the effort for archaeologists to create unrolled

painted pottery surfaces and opens up the entire set of published or available photos for further analysis. This step of creating almost distortion free images is crucial in vase painting studies. Only then, a direct analysis and reasonable comparison of motifs and several characteristic features will be achieved. Moreover, complete and distortion-minimized motif depictions can ease computer-aided tasks of pattern recognition, retrieval, and variability analysis. Since our algorithm consists of a number of independent modules that are arranged in a pipeline, individual components can be exchanged without breaking the whole system. This setup fosters tweaking of existing pipeline steps, even exchanging modules for new ones.

We imagine that based on the work presented in this paper, the construction of a fully automatic pipeline for mass processing and elastic unrolling of large image collections is possible. To this end, future work needs to address several issues of robustness, and reducing the number of parameters that are introduced throughout our system. A second direction will be the design of an enhanced visual interface that provides the user with means to intuitively and effectively design and control the generation of unrollings from her photographs.

6. Acknowledgements

We want to thank Stefan Schmidt from the Bavarian Academy of Sciences and Humanities and Head of the CVA Germany, for providing references to photographic guidelines for publishing ancient vases within the international framework of the CVA. We also thank Paul Bayer and Kerstin Bauer for creating new photographs of the ancient vases of the archaeological collection at the University of Graz. This work was co-funded by the Austrian Science

Fund FWF and the State of Styria, Austria within the project *Cross-modal Search and Visual Exploration of 3D Cultural Heritage Objects* (P31317-NBL).

References

- [BA83] BURT P. J., ADELSON E. H.: A multiresolution spline with application to image mosaics. *ACM Trans. Graph.* 2, 4 (Oct. 1983), 217–236. [3](#), [5](#)
- [BID*17] BANTERLE F., ITKIN B., DELLEPIANE M., WOLF L., CALIERI M., DERSHOWITZ N., SCOPIGNO R.: Vasesketch: Automatic 3d representation of pottery from paper catalog drawings. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)* (2017), vol. 01, pp. 683–690. [3](#)
- [BKMK10] BECHTOLD S., KRÖMKER S., MARA H., KRATZMÜLLER B.: Rollouts of fine ware pottery using high resolution 3d meshes. In *Proceedings of the 11th International Symposium on Virtual Reality, Archaeology and Cultural Heritage* (Goslar, 2010), Artusi A., Joly M., Lucet G., Pitzalis D., Ribes A., (Eds.), VAST '10, Eurographics Association, pp. 79–86. [2](#)
- [BL06] BROWN M. A., LOWE D.: Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision* 74 (2006), 59–73. [3](#)
- [Blo57] BLOESCH H.: Conseils pour photographier les vases grecs. In *Colloque international sur le Corpus vasorum antiquorum, Lyon, 3–5 juillet 1956*, Dugas C., Metzger H., (Eds.), Centre national de la recherche scientifique, Lyon, 1957, pp. 31–33. [2](#), [4](#)
- [DTN*20] DUNG H. V., TRAN D.-P., NHU N., PHAM T.-A., PHAM V.-H.: *Deep Feature Extraction for Panoramic Image Stitching*. 03 2020, pp. 141–151. [3](#)
- [FB81] FISCHLER M. A., BOLLES R. C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24, 6 (June 1981), 381–395. [2](#)
- [KBMM19] KARL S., BAYER P., MÁRTON A., MARA H.: Advanced documentation methods in studying corinthian black-figure vase painting. In *Digitize: Research – Record – Reactivate – Reconstruct, Proceedings of the 23rd International Conference on Cultural Heritage and New Technologies 2018, Vienna, November 12.–15. 2018* (Vienna, 2019), Börner W., Uhlirz S., (Eds.), CHNT '19, Museen der Stadt Wien - Stadtarchäologie, pp. 1–13. [2](#)
- [KPP96] KARRAS G., PATIAS P., PETSA E.: Digital monoplotting and photo-unwrapping of developable surfaces in architectural photogrammetry. *International Archives of photogrammetry and Remote Sensing* 31 (1996), 290–294. [4](#)
- [LDW58] LENORMANT C., DE WITTE J.: *Élite des monuments céramographiques: matériaux pour l'histoire des religions et des moeurs de l'antiquité*, vol. 3. Leleux, Librairie-Éditeur, Paris, 1858. [2](#)
- [LKL*20] LENGAUER S., KOMAR A., LABRADA A., KARL S., TRINKL E., PREINER R., BUSTOS B., SCHRECK T.: A sketch-aided retrieval approach for incomplete 3d objects. *Computers & Graphics* 87 (2020), 111–122. [3](#)
- [Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *International journal of computer vision* 60, 2 (2004), 91–110. [2](#), [5](#)
- [LWL*18] LI J., WANG Z., LAI S., ZHAI Y., ZHANG M.: Parallax-tolerant image stitching based on robust elastic warping. *IEEE Transactions on Multimedia* 20, 7 (2018), 1672–1687. [3](#), [5](#), [6](#)
- [Mom02] MOMMSEN H.: Zur Auswertung von Formfotos und Profilzeichnungen. In *Vasenforschung und Corpus Vasorum Antiquorum: Standortbestimmung und Perspektiven*, Bentz M., (Ed.), Corpus Vasorum Antiquorum, Deutschland, Beiheft 1. Bavarian Academy of Sciences and Humanities, C.H.Beck, Munich, 2002, pp. 23–36. [2](#), [4](#)
- [MP13] MARA H., PORTL J.: Acquisition and documentation of vessels using high-resolution 3d-scanners. In *Interdisziplinäre Dokumentations- und Visualisierungsmethoden*, Trinkl E., (Ed.), Corpus Vasorum Antiquorum, Österreich, Beiheft 1. Austrian Academy of Sciences, 2013, pp. 25–40. [2](#)
- [Ots79] OTSU N.: A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics* 9, 1 (1979), 62–66. [6](#)
- [PKBS18] PREINER R., KARL S., BAYER P., SCHRECK T.: Elastic Flattening of Painted Pottery Surfaces. In *Eurographics Workshop on Graphics and Cultural Heritage* (2018), Sablatnig R., Wimmer M., (Eds.), The Eurographics Association. [2](#), [6](#), [9](#)
- [Ric15] RICE P. M.: *Pottery Analysis. A Sourcebook*, 2 ed. The University of Chicago Press, Chicago and London, 2015. [1](#)
- [RMK13] RIECK B., MARA H., KRÖMKER S.: Unwrapping highly-detailed 3d meshes of rotationally symmetric man-made objects. In *Proceedings of the XXIV International CIPA Symposium, 2–6 September 2013, Strasbourg* (2013), ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences, vol. II-5/W1, pp. 259–264. [2](#)
- [S*85] SUZUKI S., ET AL.: Topological structural analysis of digitized binary images by border following. *Computer vision, graphics, and image processing* 30, 1 (1985), 32–46. [4](#)
- [Sap18] SAPIRSTEIN P.: A high-precision photogrammetric recording system for small artifacts. *Journal of Cultural Heritage* 31 (06 2018), 33–45. [2](#)
- [Sch14] SCHMIDT S.: Cva today. some remarks. In *Les cratères à volutes, Actes du Colloque de l'Académie des Inscriptions et Belles-Lettres, Paris, 8 juin 2012*, De La Genière J., (Ed.), Académie des Inscriptions et Belles-Lettres, Paris, 2014, pp. 9–12. [4](#)
- [Sti02] STISSI V. V.: *Pottery to the people. The production, distribution and consumption of decorated pottery in the Greek world in the Archaic period (650-480 BC)*. PhD thesis, University Amsterdam, 2002. [1](#)
- [Vil65] VILLARD F.: *Corpus Vasorum Antiquorum*, Musée du Louvre 13, France 21. Librairie Ancienne Honoré Champion, Paris, 1965. [2](#)
- [Wal08] WALTER C.: Towards a more 'scientific' archaeological tool. the accurate drawing of greek vases between the end of the nineteenth and the first half of the twentieth centuries. In *Archives, Ancestors, Practices: Archaeology in the Light of its History*, Schlanger N., Nordbladh J., (Eds.), Berghahn Books, New York, 2008, pp. 179–190. [2](#)
- [ZE18] ZIMMERMANN-ELSEIFY N.: *Attisch schwarzfigurige Lekythen*. *Corpus Vasorum Antiquorum*, Berlin, Antikensammlung 17, Deutschland 102. Bavarian Academy of Sciences and Humanities, C.H.Beck, Munich, 2018. [9](#)