

Context-based Surface Pattern Completion of Ancient Pottery

Stefan Lengauer¹ , Reinhold Preiner¹ , Ivan Sipiran² , Stephan Karl³ , Elisabeth Tinkl³ ,
Benjamin Bustos² , and Tobias Schreck¹ 

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

²University of Chile, Department of Computer Science, Chile

³University of Graz, Institute of Classics, Austria

Abstract

Among various ancient cultures it was common practice to adorn pottery artifacts with lavish surface decoration. While the applied painting styles, color schemes and displayed mythological content may vary greatly, the presence of simple patterns which appear in a repetitive manner can be observed across civilizations and periods. Such pattern sequences generally are arranged in a structured manner in ornament bands or columns that extend over the entire surface of the object. Due to the poor conservation state of many cultural heritage objects, parts of the surface are oftentimes badly damaged or missing altogether. Yet, if the majority of a pattern sequence is preserved, this information can be leveraged to approximate its missing parts. We present an approach that allows the fully automatic determination of the generation rule inherent to a repetitive surface pattern. Based on this generation rule and the preserved patterns from the same pattern class we propose a workflow for reconstruct missing or damaged parts of the surface painting. We evaluate our approach by applying it to a selection of pottery from ancient Peruvian and Greek cultures, showing that our automatic approach is able to handle a variety of problem cases.

CCS Concepts

• **Applied computing** → Arts and humanities; • **Information systems** → Information retrieval;

1. Introduction

The digital restoration of cultural heritage (CH) artifacts is a well know research objective, as the physical restoration is usually not possible due to various reasons, *e.g.*, cost factors, the fragility of objects, or different parts of one and the same object being curated by different museums. Yet, with the advancement of 3D scanning hardware more and more objects become available as detailed 3D models, allowing the application of computer-based restoration techniques. Papaioannou *et al.* [PSA*17] differentiate thereby between *reassembly*, the identification of affiliations between fractured artifact parts and their realignment, and *object completion*, the approximation of objects parts which are missing due to erosion or an incomplete set of object fragments. Our proposed approach belongs to the latter, but in contrast to the majority of techniques which deal with the completion of an object's geometry [SGS14, GSP*14, APM15, Sip17, Sip18, HS18, LBB19, SAC004, HTG14a, HTG14b], we focus on the completion of its texture. This is particularly relevant for the analysis of ancient pottery artifacts which oftentimes exhibit lavish surface decorations. The painting style, color scheme and illustrations vary greatly between different cultures and periods, but an almost ubiquitous theme are sequences of repetitive patterns. Those are usually arranged in rows or columns around the objects' solid of revolution. Due to the various stages of degradation, characteristic for ancient pottery, elements of a sequence are oftentimes eroded beyond recognition or

missing altogether (see examples in Fig. 1). However, by exploiting concepts of regularity and similarity, it is possible to reconstruct missing or eroded parts.



Figure 1: A selection of pottery from the Josefina Ramos de Cox museum, Lima, exhibiting various stages of erosion.

With our proposed approach we aim to complete missing patterns in incomplete ornamentation sequences. Apart from handles and other appendages, pottery objects are generally very regular solids of revolution due to their method of production on a potter's wheel. This allows us to deal with the completion problem in 2D, as such 3D surfaces can be efficiently mapped into image space [RMK13, PKBS18, KPP96, EGAND06, CRG*13]. Starting from the projected planar surface, our approach comprises two main steps. First, the detection of sequence gaps for a given sur-



Figure 2: The original surface with the border of the inpainting mask in blue (left), and the patch-inpainted [CPT04] result (right).

face pattern (Sec. 4); and second, the filling of the gap with meaningful texture (Sec. 5). While the completion of missing image parts is a well known image processing task, known as *inpainting* [JBAA*19, WWH*21, AAS21, CPT04], such approaches are generally agnostic *w.r.t.* global pattern arrangements and are not aware of any missing patterns given an empty space on the surface to fill. In Fig. 2 we provide an example of the outcome of a the patch-inpainting approach by Criminisi *et al.* [CPT04], applied to the surface of the pre-Columbian pot shown in Fig. 1, right. It can be seen that, while the generated texture at the vicinity of the mask borders appears to be indistinguishable from the original texture, the result gradually worsens towards the center of the inpainted region where the generated texture is a mix of several unrelated image regions. This is to be expected as general purpose inpainting algorithms are not able to capture the essence of the sequence, but merely reconstruct the target area based on the border gradients and colors. We conclude that off-the-shelf inpainting algorithms are generally not able to faithfully reconstruct ornamentation sequences.

In contrast, we employ a workflow where we estimate the positions, orientations and scales of missing patterns based on the sequence's generation grammar, before we employ a statistical denoising of the respective areas based on the remaining elements of the sequence. Note that our proposed approach focuses exclusively on the reconstruction of textual information instead of geometric information. The pipeline is discussed in depth in Sec. 3, before the application of this workflow on a selection of real world CH objects is presented in Sec. 6. As an additional validation we also experiment with a synthetic surface abrasion for objects which do not exhibit surface deficiencies. This has the advantage that we can increase the number of valid reconstruction targets while also having a ground truth for the completion process. In Sec. 7 we discuss limitations and possible targets for future work before we conclude the paper with a discussion of possible applications in the CH domain in Sec. 8.

2. Related Work

Restoration of defects in digital 3D cultural heritage objects has several aspects, and requires solving different tasks. That is, it has been an intensive research focus, especially over the last decade. First attempts to tackle the problem of reproducing surface patterns in CH objects were mainly focused in geometric methods with user intervention. Kolomenkin *et al.* [KLST11] proposed a framework to recreate surface relief patterns from line drawings entered by a user. The method optimizes the 3D relief geometry for a given drawing by imposing smoothness constraints characterized by the Laplacian of the relief. These constraints enhance the smooth tran-

sitions between the reliefs and give a realistic appearance of the reconstructions.

Andreadis *et al.* [APM15] propose a semi-automatic pipeline for the geometric reassembly of fractured 3D objects. The pipeline consists of several steps. First, a preprocessing step produces a segmentation of the object surfaces, and classifies each segment as either 'intact' or 'fractured'. Second, the method matches the fragments according to their fractured surfaces. Third, the method performs a pairwise surface registration for fragments exhibiting significant erosion or a missing part. Fourth, the object is reassembled using a graph-based approach. Finally, a last step, symmetry-based fragment registration is used to align any remaining disconnected fragment. Each of these steps requires human interaction to ensure a correct reassembly process, and to help the system aligning fragments that are too eroded or have significant missing parts.

Lamb *et al.* [LBB19] propose an automatic method to reconstruct missing geometry of a 3D object. To this end, the algorithm requires a reference 3D object. In a first step the incomplete object is aligned with this reference object. Using the information from the alignment step, the exterior of the missing part is generated. Finally, the algorithm computes a smooth surface transition along the outlines of the fracture. The generated fragment can be 3D printed for restoration of the broken object. Setty and Mudanagudi [SM18] propose a similar reconstruction method of missing parts for 3D point clouds. The method relies on a set of 'exemplar' models, which are geometrically similar objects to the one that is being repaired. The method selects the regions from the exemplar objects that best fit the missing parts, and uses them for the reconstruction of the object.

In contrast to these exemplar based methods [LBB19, SM18], 'hole filling' approaches [SACO04, HTG14a, HTG14b] try to reproduce a missing part of an object's geometry by patching it with a context-based or coherent surface patch obtained from the input object itself. Sharf *et al.* [SACO04] propose an automatic method to fill holes or to complete point-sampled surface areas in 3D models. The main idea of this method is to fill the hole with a surface patch considering its context, that is, the characteristics of the surface surrounding the hole. The authors show some examples of the results obtained with their method and compare it to patching the hole with a smooth surface. They also show that their method is robust to noise. A disadvantage of this method is that no overall context is considered, but holes are just filled such that the patch is merely 'similar' to its surroundings. Harary *et al.* [HTG14a] propose a method for the same problem, with the difference being that it is not only context-based, but also 'coherent'. The coherence restriction imposes that the local patch must be similar to some local patch in the model. This helps avoiding some errors that can be produced by context-based methods, where the patched surface may be similar to its surroundings, but significantly different to a corresponding local neighborhood in the surface (for example, a symmetric counterpart of the patched surface). In addition, orthogonal, method by Harary *et al.* [HTG14b] follows a semi-automatic concept, and requires the user to input four points that are used to compute a curve. This curve defines 'triangle strips', which are considered constraints of the problem and aid the reconstruction process. That is, if the surface hole covers two different segments




Figure 3: Pattern completion pipeline. The input model surface is unwrapped and subject to a self-similarity detection of individual pattern entities belonging to different classes. Individual sequences of repetitive entities within these sets are then detected and their generative grammar estimated. These are input to a context-based reconstruction and substitution of missing or deteriorated pattern elements in gaps detected in these sequences (green bracket). The reconstructed texture is finally mapped back onto the input geometry. Here, two ornament sequences of snake patterns are completed (red bracket).

or parts of the 3D model (e.g., if it covers part of the hair and part of an ear of a bust), and the user-defined strips separates these segments, the method can produce coherent and context-based patches for each of these segments.

Papaioannou *et al.* [PSA*17] propose a complete geometry restoration workflow for diverse CH objects, comprised of reassembly and completion steps. For the prior, they leverage both, fracture surfaces as well as salient features. For the latter, they present a procedure based on generalized symmetries and a complementary part extraction. They demonstrate their pipeline on a number of varied CH objects. The publication is part of the European Union-funded PREdictive digitization, reStoration and degradation assessment of cultUral heritage objectS (PRESIOUS, <http://www.presious.eu>) project.

For wheel-made pottery, the largest group of CH objects, symmetry properties can be leveraged for completion efforts. Sipiran [Sip17] describes a method that completes missing geometry by replicating symmetric features. In particular, this method exploits the global-local behavior of a heat-based function defined on the 3D surface. The function is formally proven to be invariant to rotational symmetry and exhibits symmetric correspondences around the generator axis of a given object. The method is able to detect the symmetry axis of incomplete solids of revolution. This information is then used to replicate missing parts. In a later publication Sipiran [Sip18] explores the suitability of fast local feature detectors to detect repeatable regions along the rim of rotationally symmetric objects. Since these regions convey enough information about the symmetric structure of an object, they can be used to replicate existing geometry for completion purposes. The author also demonstrates the effectiveness of the method for CH objects.

One of the few approaches dealing with textual completion within a CH context is given by Assael *et al.* [ASS*22] who employ a deep learning architecture based on transformers, with the aim of restoring ancient Greek inscriptions. The proposed deep neural network, dubbed *Ithaca*, can predict missing characters in these ancient texts. It can also predict the geographical origin of a text among 84 different regions and its period of creation in a time span from 800 BCE to 800 AD. To this end, the network was specifically trained with data from texts in ancient Greek language.

It appears the majority of related work is focused with geometric reconstruction while publications dealing with textual completion of CH objects are rather sparse and seemingly limited to image-based techniques. For this task of image-inpainting, an abundance of approaches was proposed over the years [JKW*21]. However, there is to our knowledge no method capable of dealing with the specifics of our repetitive pattern completion problem.

3. Overview of our Approach

Figure 3 illustrates our proposed reconstruction pipeline. The input as well as the output are 3D models which differ only in their texture. All processing in-between is conducted in image space. As a first step all entities of a given pattern class, together with their scale and orientation, are determined in a *Self-Similarity Detection* step. From this set of unordered pattern entities a sequence is determined by leveraging simple heuristics, intrinsic to pottery objects. Thereafter, the transformations between consecutive entities are computed and used to formulate the sequence's generation grammar. All these steps are described in depth in Sec. 4. If the estimated transformation between two neighbors differs greatly from the expected approximated prototypic transformation, it is taken as an indication for a gap in the sequence (Sec. 5.1). After determining the locations and transformations of reconstruction candidates, a statistical model is applied to complete the remaining texture in an optimal manner, based on the other patterns of this class (Sec. 5.2).

4. Generation Rule Detection

Given an object with repetitive surface patterns, the appearance of a specific pattern entity within a sequence can be generally estimated using a generation rule. That is, we assume a Markovian process [Gag17], meaning that the location, orientation and size of a pattern entity can be fully described as a transformation applied on its predecessor. Within a sequence these transformations between consecutive patterns are quite similar, unless there is a gap. Hence, the first objective is to find the prototypic transformation which describes the transformations between the uninterrupted parts of the sequence in an optimal manner. To this end, we use a four-stage process involving: (i) the detection of the individual occurrences of

a pattern, (ii) the detection of their inherent sequence (if present), (iii) the determination of the transformations between neighbors, and (iv) the estimation of the prototypic transformation. Prior to the first stage, we project the input, a textured 3D model, to a 2D plane. This greatly simplifies all processing steps and is possible for most rotationally symmetric shapes without introducing much distortion. At the end of the pipeline the inverse projection is used to remap the reconstructed texture back onto the 3D model. Several methods for the unwrapping of 3D surfaces have been published. Specifically, for the CH domain, approaches using proxy geometries [RMK13] or mass-spring systems [PKBS18] have been proposed. In our experiments we use a variation of the cylindrical unwrapping by Karras et al. [KPP96].

The detection of the individual entities of a given surface pattern can be formulated as a self-similarity detection problem, a very well researched topic in computer vision [BSI08, CPZ09, SI07, VGVZ09, LKK*20]. Note that this part of our pipeline can be performed by any of these methods providing areas and orientations of target entities, and is thus not within the scope of this paper. In our experiments, we rely on an already annotated pottery benchmark dataset (Sec. 6). That is, we assume as input a set of unordered pattern entities $\{S_i\}_{i \in I}$, with the index set I . Each entity has to convey (at least) the attributes orientation and scale (w.r.t. a reference/pivot entity) as well as location. Yet, scale and location properties can also derived if the entities are given as areas (sets of pixels or surface parts). Based on this input we determine the presence of structured arrangements. More specifically, in the case of ancient pottery, repetitive patterns are commonly arranged in rows, or – in rare cases – also in columns. Patterns which are scattered in random fashion all over the object’s surface are beyond the scope of our approach. We estimate the locations of rows K at offsets $Y = \{y\}_{k \in K} \subset \mathbb{R}_+$ and columns L at offsets $X = \{x\}_{l \in L} \subset \mathbb{R}_+$, based on the pattern entities’ centroid locations $\mathbf{C} = \{\mathbf{C}_i\}_{i \in I} \subset \mathbb{R}^2$. To this end, a *Gaussian Mixture Model* (GMM) is fitted to their respective x and y -coordinates. The most likely number of mixtures, reflecting the number of rows or columns, is determined using the *Bayesian Information Criterion* (BIC) [Sch78] by comparing the quality of the fits with $n \in [1..n_{max}]$ components. The maximum number of rows or columns n_{max} is limited to $|I|/3$ as we require each row or column to feature at least 3 pattern entities. I_k and I_l denote the index sets of pattern entities belonging to the k -th row (sorted by their x -coordinate) and the l -th column (sorted by their y -coordinate) respectively.

The relation between the pattern S_i and its predecessor S_{i-1} can be formulated as $S_i = \hat{M}_{i-1,i} S_{i-1}$, with the most likely transformation $\hat{M}_{i-1,i} = \arg \max_{M \in \mathcal{M}} \text{sim}(S_i, M S_{i-1})$ and an arbitrary similarity function sim . The task of finding $\hat{M}_{i-1,i}$ in the set of possible transformations $\mathcal{M}$ is an image registration problem [Bro92]. For our application we assume that $\hat{M}_{i-1,i}$ can be sufficiently well described with an affine transformation. With this class of geometric transformations a combination of a linear transformation and a translation can be modelled. In the 2D domain this transformation is usually represented with a 2×3 matrix

$$M = \begin{bmatrix} c_x \cos(\phi) & -s_x \sin(\phi) & t_x \\ s_y \sin(\phi) & c_y \cos(\phi) & t_y \end{bmatrix}, \quad (1)$$

with c_x and c_y as the scaling, s_x and s_y as the shear, ϕ as the ro-

tation and t_x and t_y as the translation. For our application we assume a uniform scaling $c = c_x = c_y$ and the absence of any shear $s_x = s_y = 0$. We denote the transformation from pattern $i-1$ to pattern i with four parameters as $\theta_i = \langle c_i, \phi_i, t_{x_i}, t_{y_i} \rangle^T$. To determine the prototypic transformation between any consecutive patterns, with parameters $\mu = \langle \tilde{c}, \tilde{\phi}, \tilde{t}_x, \tilde{t}_y \rangle^T$, we fit a multivariate *Students t-distribution* [Stu08] $\text{St}(\Theta|\mu, \lambda, \mathbf{v})$ to the set of transformations $\Theta = \{\theta_i\}_{i \in I}$, with $\mathbf{v}$ as the degrees for freedom and λ as the inverse precision scaling. We selected this distribution model as it is robust w.r.t. outliers for a small sample size.

5. Content-based Reconstruction

The content-based pattern recognition is a two-stage process. First, gaps in the pattern sequence are detected (Sec. 5.1), before the number of missing patterns, their locations, orientations and visual appearance is approximated in a gap filling step (Sec. 5.2).

5.1. Gap Detection

For the identification of gaps we leverage the prototypic transformation parameters μ . We define gap candidates as the patterns whose transformation w.r.t. their predecessor, given by θ_i , differs significantly from μ . That is, the index set of gap candidates $\mathcal{G}$ is formally given by

$$\mathcal{G} = \{i \in I : \max |\theta_i - \mu| > t_g \sigma\}, \quad (2)$$

with $\sigma = \sqrt{|I|^{-1} \sum_{i \in I} (\theta_i - \mu)^2}$ as the standard deviations of the transformation parameters and t_g as a cut-off threshold. For practical applications $t_g = 2$ showed to provide satisfying results. Note, that erroneous gap candidates, resulting from a too low threshold, are usually no concern as it will be recognized in the gap filling step that no pattern will fit in-between.

5.2. Gap Filling

For each of the gap candidates $g \in \mathcal{G}$ (Eqn. (2)) we add $n_g = \max\{\theta_g \oslash \mu\} - 1$ ($\oslash$ denotes the element-wise Hadamard-division) patterns to fill the void. We determine their respective transformations such that the space between the last pattern before the gap and the next preserved pattern is optimally bridged. I.e., we expect the bounding box $bbox_j$ of the j -th new pattern to be at $bbox(S_{g-1})M(\theta_g \oslash j + 1)$, with $M(\theta)$ as the affine transformation matrix (Eqn. (1)) for parameters θ .

In order to estimate the original texture at $bbox_j$ we perform a Bayesian image denoising [KY19], which allows us to optimally infer the noise-free version $\mathbf{y}$ of a noisy input $\mathbf{x} = [x_0, x_1, \dots, x_{D-1}]^T$, with $D = w \cdot h$ and w, h as the image’s width and height. With $\mathbf{y}^* = [y_0^*, y_1^*, \dots, y_{D-1}^*]^T$ denoting the corresponding noise free image, we formulate the relation between $\mathbf{x}$ and $\mathbf{y}^*$ as $(\mathbf{x})_i = (\mathbf{y}^*)_i \xi$, $i \in [0..D-1]$, with ξ being sampled from a Gamma distribution, as we assume the noise in $\mathbf{y}$ to be multiplicative and independent. As $\mathbf{y}^*$ is unknown, the goal is to find the best fitting $\hat{\mathbf{y}}$, given the observed $\mathbf{x}$. Such can be achieved with the *Minimum Mean Squared Error* function

$$\hat{\mathbf{y}}_{MMSE}(\mathbf{x}) = \frac{\sum_{\mathbf{y}' \in \mathcal{Y}} \mathbf{y}' p(\mathbf{x}|\mathbf{y}')}{\sum_{\mathbf{y}' \in \mathcal{Y}} p(\mathbf{x}|\mathbf{y}')},$$

with $\mathbf{Y}$ as the remaining noise-free entities of the pattern class and

$$p(\mathbf{x}|\mathbf{y}') = \frac{1}{\eta^{kD}\Gamma(k)^D} \exp\left(-\frac{\lambda}{\eta} \mathbf{1}^T (\mathbf{x} \odot \mathbf{y}') + \lambda(k-1) \mathbf{1}^T \log(\mathbf{x} \odot \mathbf{y}')\right)$$

as likelihood of the multiplicative Gamma noise. k is the shape parameter and η the scale parameter of the Gamma distribution and $\lambda \in [0, 1]$ is a governing parameter for the ‘peakyness’ of the likelihood. $\lambda = 1$ results in $\mathbf{x}$ being equivalent to exactly one element of $\mathbf{Y}$, while $\lambda < 1$ results in a blend from multiple of the best matching noise free images. $\lambda = 10e-2$ was used in all experiments.

In our case the noisy input $\mathbf{y}$ is the texture observed at $bbox_j$ and the noise free references $\mathbf{Y}$ are the textures at $\{bbox(\mathcal{S}_i)\}_{i \in I}$. In order to employ the statistical model all inputs must be brought into a normalized representation, meaning that all images must have the same dimensions and display the pattern entities in the same manner. To this end, we, first of all, transform all elements to match the appearance of the pivot element. *I.e.*, for the pattern $\mathcal{S}_i$ we apply the inverted transformation of all its prior sequence elements $\prod_{k=1}^{i-1} \hat{\mathcal{M}}_{k-1,k}^{-1}$. Note that the transformations have to be applied in reverse order since affine transformations are not commutative. This aligns the bounding boxes of all patterns with the bounding box of the pivot element. Additionally, they must be resized to a uniform size of $\bar{w} \times \bar{h}$ pixels with $\bar{w} = |I|^{-1} \sum_{i \in I} width(bbox(\mathcal{S}_i))$, $\bar{h} = |I|^{-1} \sum_{i \in I} height(bbox(\mathcal{S}_i))$ as the patterns’ mean width and height respectively. Lastly, we match the patterns’ histograms with a sort matching approach [RVBA00]. This algorithm allows to match the histogram of a source image to the histogram of a target image by sorting the pixels in both images by intensity. The source image is then assigned the pixel intensities of the target image according to the sorted sequence. We adapt this algorithm, originally intended for gray images, for colored inputs by performing the sorting and assignment for each channel separately in the *CIE L*a*b** color space [C*04]. As we do not wish however to match images pairwise but bring a set of images into a normalized representation, we match each of them to the mean histogram computed over all images. The normalized reference images a staircase pattern on a pre-Columbian pot (Fig. 7) are featured in Fig. 4. The same normalization steps are also applied to the target texture at $bbox_j$. After obtaining the denoised version of the j -th new pattern, $\hat{\mathbf{y}}_j$, we retransform it back to the predicted location and merge it with the original texture using a multiresolution spline technique [BA83].

6. Results

The following results section is tripartite. In the first part (Sec. 6.1) we visually assess the results obtained with our pattern completion approach on actual eroded real-world artifacts. In the second part (Sec. 6.3) we look at artificially degenerated surface paintings and compare the results of our pipeline to the original surface. To this end, we employ a custom made *synthetic surface abrasion* process (Sec. 6.2), which aims to mimic the traits of actual wear. Finally, we also provide quantitative results (Sec. 6.4) by generating a large quantity of said synthetic abrasions and evaluate the performance of our method in the face of different degrees of degeneration by means of a quality measure.



Figure 4: A selection of reference samples of a staircase pattern on a pre-Columbian pot after alignment and histogram matching.

6.1. Qualitative Assessment

For the qualitative assessment of our method we look at two real-world objects, exhibiting characteristic weather-worn deficiencies in their surface paintings. First, a pre-Columbian pot (Fig. 5, top) from the Josefina Ramos de Cox museum in Lima, Perú. The annotation of its surface patterns is publicly available (<https://datasets.cgv.tugraz.at/pattern-benchmark/>) from a previous annotation effort [LSP*21]. This vessel features polychromatic surface paintings with two distinct pattern classes – a series of black triangular staircase patterns with alternating orientation at the top and the bottom; as well as a snake-shaped pattern appearing in three parallel rows with black and white colorization. The pattern sequence in the middle row comprises four severely damaged pattern entities. From the results of the automatic completion (Fig. 5, top right) it can be seen that the newly generated texture blends in seamlessly and is almost indistinguishable from the original surface.

The second exhibit – an Attic black-figured amphora from the Kunsthistorisches Museum in Vienna, Austria (<https://odeeg.acdh.oeaw.ac.at/vases/object/detail/127>) – stems from ancient Greek culture. Its surface patterns are annotated in a similar fashion to the pre-Columbian patterns, with the annotation given in the supplementary material. This vessel exhibits extensive surface paintings comprising several human and animal figures together with a variety of abstract background fillings. Noteworthy *w.r.t.* pattern completion is the sequence of linked bud and flower shapes on the lower body of the amphora, which is damaged in two areas (Fig. 5, bottom left). Yet, compared to the previous example, those are better preserved and in some cases there is just a fraction of the pattern missing. Fig. 5, bottom right, shows the results of the completion effort which is used to reconstruct two pattern entities in first gap and four in the second one. This example shows the benefits of the Bayesian denoising, since the reconstruction is very faithful to the orientation and overall appearance of the incomplete patterns.



Figure 5: Eroded originals (left) and our reconstruction (right) of a snake shape on a pre-Columbian pot (top, Josefinas Ramos de Cox museum, Lima) and a bud/flower shape on the lower body of an Attic black-figured amphora (bottom, Kunsthistorisches Museum, Vienna).

6.2. Synthetic Surface Abrasion

The examples in Section 6.1 are comprised of objects with real surface deficiencies, showing the applicability of our method for real-world data. On the downside, however, these do not allow to objectively evaluate the quality of the completion, since the original appearance of the surface pattern is unknown. As a consequence, we generate additional realistic examples by synthetically introducing plausible chippings to the object surfaces. The unchanged originals serve as the ground truth for the missing surface parts.

The synthetic holes are generated in a completely randomized manner on the unrolled surface, but with constraints which ensure that they resemble realistic chippings. That is, we model an arbitrary abrasion mask by starting with a circle with random radius $r \in [r_{min}, r_{max}]$ ($r_{min} = 0.01d_s$ and $r_{max} = 0.05d_s$, with d_s as the diameter of the unrolled surface), which we place at a random location on the surface. To mimic the irregular shape and jaggedness of real holes we superimpose a random noise ϵ to the radius of the circle, amounting to $\epsilon(\theta) = \sum_{e \in [0..E]} a_e / e \sin(2^e \theta + \phi_e)$ at polar angle θ , with E as the number of error functions, determining the frequency of the border noise, and $\{a_e\}$ and $\{\phi_e\}$ as random amplitudes and offsets. The number and irregularity of holes is governed




Figure 6: 300×300 pixel erosion masks, visualizing the impact of the expected rate of hole occurrence k_h and number of error functions E on the synthetic abrasion process.

by $n_h \sim \Gamma(k_h)$ and E respectively. The influence of those parameters can be seen from Fig. 6. $k_h = 5$ and $E = 15$ have been selected for our experiments.

By flood-filling the generated holes, which can of course overlap, we obtain a mask $\mathcal{M}$ which we use to selectively delete surface texture. Instead of leaving the holes as colourless void we fill them with a salient green hue which makes them easy to recognize and also allows us to investigate the influence of the texture in these areas on the reconstruction. Applying the mask as it is to the surface would result in unnaturally sharp edges at the borders of the holes. Instead we use an alpha blending in the vicinity of those edges to smoothen the transitions. To this end, we employ an alpha mask $\mathcal{M}_\alpha = 1 - \exp(-\mathcal{M}_{dist}^2 / 2\sigma_\alpha^2)$, with $\sigma_\alpha = 8$ governing the smoothness. $\mathcal{M}_{dist}$ is the mask $\mathcal{M}$ after a distance transform [Bor86], which gives at each pixel the distance to the closest zero pixel. The surface image with the synthetic abrasion $\mathcal{I}_{abr} = \lfloor \mathcal{I} \cdot (1 - \mathcal{M}_\alpha) + \mathcal{M} \cdot \mathcal{M}_\alpha \rfloor$ is obtained by combining the original surface $\mathcal{I}$ with the generated mask $\mathcal{M}$.

6.3. Qualitative Synthetic Experiments

To qualitatively assess the pattern completion after applying a synthetic abrasion, we select three distinct objects with diverse characteristics, and thus posing different challenges. They are displayed in Fig. 7, left column. The first one is a pre-Columbian bowl shape, exhibiting three rows of ‘N’-shaped patterns with white and reddish hue. The second object is a pot with two different pattern types – a very simple circular pattern as well as triangular staircase pattern, appearing with alternating orientations. The third object, an Attic Geometric high-rimmed bowl (<https://odeeg.acdh.oew.ac.at/vases/object/detail/100>) is the most complex input. Although, its surface painting is monochro-




Figure 7: Three exemplaric pottery objects with their original 3D model (first column), their model after applying the synthetic surface abrasion (second column), after the pattern completion with our method (third column) and the color-coded difference between original and reconstruction (fourth column). The artificial chip-offs are filled with uniform green color.

matic (black), it features three distinct pattern classes which appear in a repetitive manner: an organic bird shape, a cross-hatched triangle and a stack of inverted 'V's. All of them appear a total of 26 times around the solid of revolution. The second column in Figure 7 illustrates the inputs after they are subjected to the synthetic surface abrasion. Note, that the generated artificial chipings (in green color) vary greatly in shape and size. The third column shows the models from the previous column after applying the automatic pattern completion while the fourth column constitutes a quality map revealing the differences between the Ground Truth and the Reconstruction. These maps are obtained by computing the pixel-wise L_2 norm of the respective RGB color channels. For the first object, the pre-Columbian bowl, it can be observed that the reconstructed patterns in the top and bottom row are visually indistinguishable from the existing ones. A slight inaccuracy can be observed with the red patterns in the middle row, but they also appear plausible at large. Note that some of the green color, marking the surface abrasion, is still visible below the bottom row. This can be attributed to the fact that this region is outside the bounding boxes of the patterns and is thus not part of the reconstruction process. Thin outlines of the depicted patterns are perceptible in the quality map, indicating that reconstructed patterns' shapes marginally differ from the originals'. The pre-Columbian pot in the second row features a gap with seven missing patterns of alternating orientation. The completion is able to correctly predict their orientations. However, some ghosting artifacts can be observed. For the Geometric bowl, four of the bird patterns are reconstructed but no errors could be spotted from neither the visual inspection nor the quality map.

6.4. Quantitative Synthetic Experiments

Besides the qualitative assessment in Sec. 6.3, which shows just one specific surface abrasion use case, we also conduct a quantitative evaluation based on an objective metric. We leverage the quality maps (also shown in Fig. 7, right) to compute a single quality measure per experiment. To this end, we integrate over this map and divide the result by the number of pixels which are to be reconstructed and the maximum value of color intensities. This gives a value between 0 and 1, where 0 describes a perfect match and 1 indicates the complementary colorization. The same measure is also determined for the synthetically damaged texture, prior to the completion. We compute these pairs of measures for 100 different synthetic abrasion experiments for each of the three objects in Fig. 7. The resulting measurements, plotted over the relative removed surface is shown in Fig. 8. The error of the reconstruction is mostly within the range of 0.05 and 0.15 and does not seem to change with increasing amount of reconstructed surface area. The error between the ground truth and the damaged surface, however, describes a fast ascend before seemingly converges towards about 0.3. This is to be expected, as the green hue is not the ultimate complementary to the original texture.

7. Discussion, Limitations and Future Work

Apparently, at least some of the green color, marking the synthetic abrasions, is still visible after the completion process (Fig. 7). This is due to the fact that our reconstruction is limited to the areas, populated by patterns. For most practical examples this is sufficient (Fig. 5) as the clay at the chipped or worn off regions generally has a similar color than the rest of the vessel. Yet, in some cases other




Figure 8: Reconstruction error ('x') and difference between synthetically damaged surface and ground truth ('o'), over the relative removed surface area for three different input objects.

surface decoration, e.g., straight lines around the solid of revolution, which are without the scope of the proposed method, are also incomplete. For such cases a processing step combining a chipping detection [DGP*19] with an context-based inpainting [CRG*13], prior to the pattern completion, would be a viable addition to our workflow.

In some cases ghosting artifacts can be observed occasionally (see inset) at the boundaries between original and reconstructed patterns. This is especially the case if the original sequence has varying distances in-between patterns, since we assume a uniform distribution. This assumption of regularity is sometimes problematic, especially in the vicinity of an ornament band's seam, because here the ancient painter had to adjust the patterns' widths to the still available space, making them either narrower or broader. There are three ways to address this issue. First, the stitching algorithm [BA83] used to fuse the reconstructed patterns with the rest of the surface is not optimal for our workflow as it is intended for the stitching of panoramic photographs. For our application the inserted texture patches are oftentimes too small and the overlap with the rest of the surface too insignificant to generate reasonable results. However, – same as with the other parts of our pipeline – this building block can be readily replaced with an algorithm more suited for this task. Other viable remedies include a fine alignment of the patch before it is stitched or an incorporation of the user into the completion process.



Another important question is how many preserved patterns are necessary to reasonably reconstruct a sequence. Generally, the number of preserved patterns needs to be significantly higher than the number of missing ones. Our experiments indicate that a completion is feasible with up 30% of a sequence missing. With even higher fractions the generation rule detection (Sec. 4) and gap detection (Sec. 5.1) are no longer reliable.

8. Applications

The proposed approach contributes to two significant aspects of archaeological pottery research. First, for the attribution of pottery objects (so-called vases) to painters/workshops and to manufacturing techniques, and second, for educational and communicational purposes regarding these CH objects. Within the field of research on ancient vase painting the determination of a pattern sequence's generation grammar can be used for recognising characteristic properties of painters' hands or pottery workshops. For instance, friezes with rows of birds are very common in Geometric pottery. From the automatically determined generation grammar it is possible to analyse their similarity to other occurrences on different pottery objects and to cluster them accordingly. From a ceramologist's view, the estimation of the number of missing patterns is generally useful as it helps to grasp the painter's overall plan. Another important aspect of this reconstruction approach is the prediction of expected patterns, especially in cases where such repetitive patterns are only partly preserved. As a proxy, similar to the domain knowledge applied by an archaeologist processing pottery, this can strongly support computer-assisted pattern recognition or retrieval tasks. Regarding the educational field, the reconstruction of missing parts is relevant for presenting CH objects – physically as well as virtually; e.g., for teaching activities or for replicas.

9. Conclusion

We present an automatic approach for the completion of pattern sequences on ancient pottery. To this end, we determine the sequence's grammar by inferring a prototypic transformation between consecutive patterns. This generation grammar allows us to detect gaps in the sequence which will fill with a statistical denoising approach, using the sequence's preserved patterns. The employment of this workflow for real world examples, as well as synthetically generated use cases, proves the applicability of our approach for the problem at hand. Hence, the contribution of our paper is the proposed workflow with some self-created and some preexisting building blocks. We want to stress that those could be exchanged and adopted independently to improve certain aspects or address additional challenges.

Acknowledgements

This work was co-funded by the Austrian Science Fund FWF and the State of Styria, Austria within the project *Crossmodal Search and Visual Exploration of 3D Cultural Heritage Objects* (CrossSAVE-CH, P31317-NBL).

References

- [AAS21] AHMED H. O., ALFAQHERI T., SADKA A. H.: Digital image inpainting techniques for cultural heritage preservation and restoration. In *Data Analytics for Cultural Heritage*. Springer, 2021, pp. 91–122. 2
- [APM15] ANDREADIS A., PAPAIOANNOU G., MAVRIDIS P.: Generalized digital reassembly using geometric registration. In *2015 Digital Heritage* (2015), vol. 2, IEEE, pp. 549–556. 1, 2
- [ASS*22] ASSAEL Y., SOMMERSCHIED T., SHILLINGFORD B., BORDBAR M., PAVLOPOULOS J., CHATZIPANAGIOTOU M., ANDROUTSOPOULOS I., PRAG J., DE FREITAS N.: Restoring and attributing ancient texts using deep neural networks. *Nature* 603, 7900 (2022), 280–283. 3

- [BA83] BURT P. J., ADELSON E. H.: A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics (TOG)* 2, 4 (1983), 217–236. 5, 8
- [Bor86] BORGES G.: Distance transformations in digital images. *Computer vision, graphics, and image processing* 34, 3 (1986), 344–371. 6
- [Bro92] BROWN L. G.: A survey of image registration techniques. *ACM computing surveys (CSUR)* 24, 4 (1992), 325–376. 4
- [BSI08] BOIMAN O., SHECHTMAN E., IRANI M.: In defense of nearest-neighbor based image classification. In *2008 IEEE conference on computer vision and pattern recognition* (2008), IEEE, pp. 1–8. 4
- [C*04] CONSORTIUM I. C., ET AL.: Specification icc. 1: 2004-10 (profile version 4.2. 0.0) image technology colour management. architecture, profile format, and data structure, 2004. 5
- [CPT04] CRIMINISI A., PÉREZ P., TOYAMA K.: Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on image processing* 13, 9 (2004), 1200–1212. 2
- [CPZ09] CHATFIELD K., PHILBIN J., ZISSERMAN A.: Efficient retrieval of deformable shape classes using local self-similarities. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops* (2009), IEEE, pp. 264–271. 4
- [CRG*13] CORNELIS B., RUŽIĆ T., GEZELS E., DOOMS A., PIŽURICA A., PLATIŠA L., CORNELIS J., MARTENS M., DE MEY M., DAUBECHIES I.: Crack detection and inpainting for virtual restoration of paintings: The case of the ghent altarpiece. *Signal Processing* 93, 3 (2013), 605–619. 1, 8
- [DGP*19] DULECHA T., GIACHETTI A., PINTUS R., CIORTAN I.-M., JASPE VILLANUEVA A., GOBBETTI E.: Crack detection in single-and multi-light images of painted surfaces using convolutional neural networks. In *Eurographics Workshop on Graphics and Cultural Heritage* (2019), The Eurographics Association, pp. 43–50. 8
- [EGAND06] EL GLALY Y., ATTA E., NAZMY T., DESOUKY B.: Digital inpainting for cultural heritage preservation. In *International Conference on Computer Theory and Application, ICCTA* (2006), pp. 330–333. 1
- [Gag17] GAGNIUC P. A.: *Markov chains: from theory to implementation and experimentation*. John Wiley & Sons, 2017. 3
- [GSP*14] GREGOR R., SIPIRAN I., PAPAIOANNOU G., SCHRECK T., ANDREADIS A., MAVRIDIS P.: Towards automated 3d reconstruction of defective cultural heritage objects. In *GCH* (2014), pp. 135–144. 1
- [HS18] HERMOZA R., SIPIRAN I.: 3d reconstruction of incomplete archaeological objects using a generative adversarial network. In *Proceedings of Computer Graphics International 2018* (New York, NY, USA, 2018), CGI 2018, Association for Computing Machinery, p. 5–11. 1
- [HTG14a] HARARY G., TAL A., GRINSPUN E.: Context-based coherent surface completion. *ACM Transactions on Graphics (TOG)* 33 (2014), 1–12. 1, 2
- [HTG14b] HARARY G., TAL A., GRINSPUN E.: Feature-preserving surface completion using four points. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 45–54. 1, 2
- [JBAA*19] JBOOR N. H., BELHI A., AL-ALI A. K., BOURAS A., JAOUA A.: Towards an inpainting framework for visual cultural heritage. In *2019 IEEE Jordan International Joint Conference on Electrical Engineering and Information Technology (JEEIT)* (2019), IEEE, pp. 602–607. 2
- [JKW*21] JAM J., KENDRICK C., WALKER K., DROUARD V., HSU J. G.-S., YAP M. H.: A comprehensive review of past and present image inpainting methods. *Computer vision and image understanding* 203 (2021), 103147. 3
- [KLST11] KOLOMENKIN M., LEIFMAN G., SHIMSHONI I., TAL A.: Reconstruction of relief objects from line drawings. In *CVPR 2011* (2011), pp. 993–1000. 2
- [KPP96] KARRAS G., PATIAS P., PETS A. E.: Digital monoplott and photo-unwrapping of developable surfaces in architectural photogrammetry. *International Archives of photogrammetry and Remote Sensing* 31 (1996), 290–294. 1, 4
- [KY19] KATAOKA S., YASUDA M.: Bayesian image denoising with multiple noisy images. *The Review of Socionetwork Strategies* 13, 2 (2019), 267–280. 4
- [LBB19] LAMB N., BANERJEE S., BANERJEE N. K.: Automated reconstruction of smoothly joining 3d printed restorations to fix broken objects. In *Proceedings of the ACM Symposium on Computational Fabrication* (2019), pp. 1–12. 1, 2
- [LKK*20] LENGAUER S., KOMAR A., KARL S., TRINKL E., SIPIRAN I., SCHRECK T., PREINER R.: Semi-automated annotation of repetitive ornaments on 3d painted pottery surfaces. In *Eurographics Workshop on Graphics and Cultural Heritage* (2020), The Eurographics Association, pp. 1–4. 4
- [LSP*21] LENGAUER S., SIPIRAN I., PREINER R., SCHRECK T., BUS-TOS B.: A Benchmark Dataset for Repetitive Pattern Recognition on Textured 3D Surfaces. *Computer Graphics Forum* (2021). 5
- [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, pp. 165–168. 1, 4
- [PSA*17] PAPAIOANNOU G., SCHRECK T., ANDREADIS A., MAVRIDIS P., GREGOR R., SIPIRAN I., VARDIS K.: From re-assembly to object completion: A complete systems pipeline. *Journal on Computing and Cultural Heritage (JOCCH)* 10, 2 (2017), 1–22. 1, 3
- [RMK13] RIECK B., MARA H., KRÖMKER S.: Unwrapping highly-detailed 3d meshes of rotationally symmetric man-made objects. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences* (2013), 259–264. 1, 4
- [RVBA00] ROLLAND J. P., VO V., BLOSS B., ABBEY C. K.: Fast algorithms for histogram matching: Application to texture synthesis. *Journal of Electronic Imaging* 9, 1 (2000), 39–45. 5
- [SACO04] SHARF A., ALEXA M., COHEN-OR D.: Context-based surface completion. In *ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), SIGGRAPH '04, Association for Computing Machinery, p. 878–887. 1, 2
- [Sch78] SCHWARZ G.: Estimating the dimension of a model. *The annals of statistics* 6, 2 (1978), 461–464. 4
- [SGS14] SIPIRAN I., GREGOR R., SCHRECK T.: Approximate symmetry detection in partial 3d meshes. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 131–140. 1
- [SI07] SHECHTMAN E., IRANI M.: Matching local self-similarities across images and videos. In *2007 IEEE Conference on Computer Vision and Pattern Recognition* (2007), IEEE, pp. 1–8. 4
- [Sip17] SIPIRAN I.: Analysis of partial axial symmetry on 3d surfaces and its application in the restoration of cultural heritage objects. In *Proceedings of the IEEE International Conference on Computer Vision Workshops* (2017), pp. 2925–2933. 1, 3
- [Sip18] SIPIRAN I.: Completion of cultural heritage objects with rotational symmetry. In *Proceedings of the 11th Eurographics Workshop on 3D Object Retrieval* (2018), pp. 87–93. 1, 3
- [SM18] SETTY S., MUDENAGUDI U.: Region of interest-based 3d inpainting of cultural heritage artifacts. *Journal on Computing and Cultural Heritage (JOCCH)* 11, 2 (2018), 1–21. 2
- [Stu08] STUDENT: The probable error of a mean. *Biometrika* (1908), 1–25. 4
- [VGVZ09] VEDALDI A., GULSHAN V., VARMA M., ZISSERMAN A.: Multiple kernels for object detection. In *2009 IEEE 12th international conference on computer vision* (2009), IEEE, pp. 606–613. 4
- [WWH*21] WANG N., WANG W., HU W., FENSTER A., LI S.: Thangka mural inpainting based on multi-scale adaptive partial convolution and stroke-like mask. *IEEE Transactions on Image Processing* 30 (2021), 3720–3733. 2