






A Benchmark Dataset for Repetitive Pattern Recognition on Textured 3D Surfaces

Stefan Lengauer¹ , Ivan Sipiran² , Reinhold Preiner¹ , Tobias Schreck¹  and Benjamin Bustos² 

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

²Department of Computer Science, University of Chile, Chile



Figure 1: A selection of 3D ancient pottery models with varying geometric shapes. The annotated classes of surface patterns (squares) exhibit different shapes, styles while their individual entities (colored polygons) show vastly different levels of repetitiveness.

Abstract

In digital archaeology, a large research area is concerned with the computer-aided analysis of 3D captured ancient pottery objects. A key aspect thereby is the analysis of motifs and patterns that were painted on these objects' surfaces. In particular, the automatic identification and segmentation of repetitive patterns is an important task serving different applications such as documentation, analysis and retrieval. Such patterns typically contain distinctive geometric features and often appear in repetitive ornaments or friezes, thus exhibiting a significant amount of symmetry and structure. At the same time, they can occur at varying sizes, orientations and irregular placements, posing a particular challenge for the detection of similarities. A key prerequisite to develop and evaluate new detection approaches for such repetitive patterns is the availability of an expressive dataset of 3D models, defining ground truth sets of similar patterns occurring on their surfaces. Unfortunately, such a dataset has not been available so far for this particular problem. We present an annotated dataset of 82 different 3D models of painted ancient Peruvian vessels, exhibiting different levels of repetitiveness in their surface patterns. To serve the evaluation of detection techniques of similar patterns, our dataset was labeled by archaeologists who identified clearly definable pattern classes. Those given, we manually annotated their respective occurrences on the mesh surfaces. Along with the data, we introduce an evaluation benchmark that can rank different recognition techniques for repetitive patterns based on the mean average precision of correctly segmented 3D mesh faces. An evaluation of different incremental sampling-based detection approaches, as well as a domain specific technique, demonstrates the applicability of our benchmark. With this benchmark we especially want to address the geometry processing community, and expect it will induce novel approaches for pattern analysis based on geometric reasoning like 2D shape and symmetry analysis. This can enable novel research approaches in the Digital Humanities and related fields, based on digitized 3D Cultural Heritage artifacts. Alongside the source code for our evaluation scripts we provide our annotation tools for the public to extend the benchmark and further increase its variety.

CCS Concepts

• **Information systems** → Evaluation of retrieval results; Retrieval efficiency; • **General and reference** → Metrics; • **Computing methodologies** → Image processing;

1. Introduction

In the Digital Humanities, the computer aided analyses of cultural heritage artifacts receives increased attention for the tasks of comparison, retrieval and understanding of objects, their meaning and historical context. An important source of information for these analysis tasks is given by the surface texture of these objects, such as the motifs and patterns painted on ancient pottery. These often exhibit distinct geometric features, and encode valuable semantic and contextual information for the archaeological analysis. To this end, digitized 3D models that capture both the object's shape as well as their surface texture constitute a valuable basis for different computer-aided analysis tasks.

A particular problem for the domain analysis of such objects is the recognition of different repetitive patterns such as decorative ornaments, which can exhibit distinctive geometric shapes of varying styles, sizes and patterns of reoccurrence (Fig. 1). An automatic identification of similar texture patterns on such vessel surfaces can serve different important domain applications, such as digital motif reconstruction, segmentation, or similarity-based retrieval. We believe that future developments of further improved techniques tackling this particular problem would be supported or even catalyzed by the availability of an expressive reference dataset as well as an established evaluation method to compare future approaches.

In this paper, we present an expert-annotated benchmark dataset for repetitive pattern detection algorithms to address these needs. We provide an annotated dataset of different 3D models of ancient painted Peruvian vessels, exhibiting varying levels of repetitiveness in their patterns. The dataset is complemented by manual expert annotations defining a reference set for pattern classes and their similarity instances on the vessel surfaces. Furthermore, we introduce an evaluation approach that can rank different recognition techniques for such repetitive patterns. With this dataset and its accompanying evaluation benchmark, we induce and inspire the development of novel approaches for the detection of similar surface patterns, not only in the classic pattern recognition community, but also within the field of geometry processing and shape analysis, which we believe can both provide significant contributions to this research problem in the future. Our benchmark can be utilized for the evaluation of retrieval algorithms based on textured 3D models, but might also be used to test novel analysis and visual exploration techniques for surface patterns with respect to their variety, variability, and similarities across different objects. Moreover, we envision novel approaches on multi-modal shape and texture reconstruction for digital restoration of Cultural heritage artifacts based on pattern similarities. In summary, our contributions are:

1. A dataset of 82 textured meshes with repetitive surface patterns (Sec. 3), annotated by a total of 2529 reference pattern occurrences defining 102 distinct similarity classes.
2. An annotation toolset for defining such similar pattern classes for textured 3D meshes (Sec. 4).
3. Two orthogonal benchmark evaluation methods for ranking different recognition techniques based on a suitable precision measure (Sec. 6).
4. An evaluation of an incremental set of naïve sampling-based pattern detection techniques as well as a domain-related detection technique tailored to repetitive ornaments (Sec. 7).

The annotated dataset is publicly available at <https://datasets.cgv.tugraz.at/pattern-benchmark/> under the *Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International* license. The site also provides the source code of the tools used for obtaining the annotations and scripts for conducting the evaluation.

2. Related Work

The problem of finding similar patterns on 3D shapes has been the focus of several recent tracks of the long-standing Shape Retrieval Contest (SHREC, <https://www.shrec.net>), which is a series of contests for evaluating methods and algorithms related to 3D object retrieval. In the following, we will scrutinize those tracks and investigate the similarities and dissimilarities with benchmark dataset. Aside from the SHREC, many approaches have emerged, relying on engineered as well as learned features. We review those in the second part of this section.

2.1. Shape Retrieval Contest (SHREC)

The SHREC 2018 track “Recognition of geometric patterns over 3D models” [BMTB*18] proposed a benchmark for recognizing relief patterns on 3D shapes. To this end, the track organizers built a dataset using 3D-scanned archaeological artifacts. The task was to detect if an object contained none, one or several of the given geometrical patterns. Although several groups registered for submitting result to this track, the task proved to be difficult for the state-of-the-art techniques, as no submissions provided fully satisfactory results. The main difficulties reported by the participants can be summed up as: (i) dealing with real-world, noisy 3D data, (ii) handling high-resolution data, and (iii) the absence of an adequate training dataset. The results of this SHREC track shows that the general problem of finding patterns in real archaeological artifacts is a very hard problem to tackle, stressing the need for new datasets that can be used for testing and training new retrieval methods. Compared to this track our dataset provides only few geometric patterns but mostly textual patterns which furthermore appear in a repetitive fashion.

Another related SHREC 2018 track was the “Retrieval of gray patterns depicted on 3D models” [MTTW*18]. The purpose of this track was to evaluate algorithms for retrieving objects with a given painted texture pattern from a synthetically generated dataset. From six registered groups, only three submitted results for the track, revealing the high complexity of the posed challenge. All submitted results used some kind of feature vectors. Even though some of the methods achieved good effectiveness scores, the task proved to be challenging. The track's dataset has two distinct differences with respect to our dataset. Firstly, the models were generated synthetically and do consequently not exhibit the deficiencies intrinsic to captured real-world artifacts, which pose a nontrivial challenge to recognition approaches. Secondly, the model surfaces feature only binary patterns while our dataset's textual patterns are polychromatic in general.

Two years later, the SHREC 2020 track “Retrieval of digital surfaces with similar geometric reliefs” [TBG*20] proposed a contest on the retrieval of 3D shapes with similar surface patches. The organizers of this SHREC track generated a dataset of 3D shapes

with 11 different geometric reliefs. Similarly to [MTTW*18] the difference to our dataset is that the models are generated synthetically. Moreover, the patterns are purely geometric and are applied to the whole surface, while the patterns and motifs in our dataset can appear anywhere on the model. The task was to detect groups of objects that have the same geometric relief, regardless of the shapes of the objects. To ascertain this premise, the global shape of the 3D models was not relevant for comparing the objects, but the geometric relief was the only discriminating feature. Several groups participated in this task, using a wide variety of retrieval techniques including classic feature vectors and deep learning techniques. Overall, methods based on deep learning techniques obtained the best results on this track, although some of the methods based on feature vectors also achieved high effectiveness scores. This topic is related to the problem of finding geometric patterns in surfaces [GCO06], which has been already studied in the context of cultural heritage [IT11]. More specific approaches take advantage of the symmetric structure of regular patterns in geometric surfaces to detect self-similarities [MBB10, BWM*11, KBW*12, HGM14].

2.2. Pattern Detection

Thompson et al. [TB18] tackled the problem of finding color patterns over a 3D surface. For this purpose, they proposed the so-called *edge Local Binary Pattern* (edgeLBP), an extension of the Local Binary Pattern (LBP) descriptor. The idea is to characterize colorimetric patterns on the surface of the shape, which can then be used for searching and classifying objects with similar patterns. The analysis of the performance of the edgeLBP descriptor shows that it achieves a high effectiveness score, and that it is robust to changes on the surface tessellation.

Additionally, Thompson et al. [TBDC19] proposed another extension of the LBP descriptor, the so-called *Mean Point Local Binary Patterns* (mpLBP). This descriptor was designed for the retrieval and classification of objects with similar geometrical patterns over the surface of a 3D shape. The definition of pattern in this work refers to elements that are repetitive. Thus, it does not characterize local decorative patterns, but the whole object. Also, the considered patterns could be geometric (e.g., variations on the shape) or colorimetric (e.g., a painting). The paper experimentally shows that the mpLBP is effective for the retrieval of relief patterns, and that is computationally less expensive than related state-of-the-art techniques like the edgeLBP.

Lengauer et al. [LKK*20] propose a method for detecting repetitive textual ornaments appearing in friezes and ornament bands on ancient Greek pottery. To this end, they leverage the assumption that repetitive ornaments occur at equal height along the axis of rotation. The approach seemingly works well for the problem at hand, but is limited to artifacts with bichrome surface colorization.

Bogacz and Mara [BM17] developed GigaMesh, a software framework for visualizing 3D-scanned cuneiform tablets. The software uses the so-called Multi-Scale Integral Invariant filtering for describing the 3D data of the scanned tablets. From this input, it produces visualizations of the documents that improve their readability. Furthermore, it supports finding repetitive cuneiform patterns in documents, applying machine-learning methods for feature

computation. The authors state that the GigaMesh framework could be extended for related problems, like studying stone wall carvings or analyzing patterns on old documents.

Úbeda et al. [ÚSN*20] proposed a deep learning approach for spotting patterns in historical documents. Their approach works in two stages. The first stage process and index the documents from the collection, computing deep feature vectors of local regions at multiple scales using the RetinaNet network. The second stage processes and locates the patterns on the documents relevant to a query image (query-by-example) by similarity search. The authors perform an experimental evaluation using a dataset of hand-written medieval documents, showing that their approach obtains better effectiveness than the state of the art for the pattern localization task.

3. Dataset

The dataset consists of 82 textured 3D triangle meshes captured in the Josefina Ramos de Cox museum in Lima, Perú. Fig. 2 shows the collection used in our dataset. It is a subset of a collection of 963 objects scanned by a structured-light scanner. The collection comprises several shape classes, such as jars, pitchers, bowls, figurines, basins, pots, plates, and vases. Moreover, the objects are attributed to several pre-Columbian cultures, each with their own characteristic artistic styles in shapes and paintings. After the scanning, the models were manually normalized to align the rotation axis with the Y+ axis and rescale its bounding box to a unit cube. The pre-processed dataset has about 130K triangles per model.



Figure 2: Textured 3D models of the dataset.

The original physical vessels also exhibit different impairments visible in the captured 3D models, such as cracks and holes in the material (Fig. 3 top left), but also chipped off and worn off surface areas (Fig. 3 bottom left and bottom right respectively) that deteriorate the final texture. Additionally, some objects exhibit missing surface parts due to occlusions and other errors in the scanning process (Fig. 3 top right). These imperfections characteristic for captured real-world models add another quality level to this dataset, and pose an additional challenge to algorithms operating on them.

4. Annotation

Any of the selected objects features between one and three distinguishable similarity classes of surface patterns (henceforth referred to as *pattern archetype*), each of which have between 2 and 134 (mean = 24.8) separable occurrences (henceforth referred to



Figure 3: Impairments of the data incorporate both, geometric deficiencies like missing parts due to decay (top left) and scanning errors (top right), as well as textual deficiencies due to chipped off (bottom left) or worn off (bottom right) surface parts.

as *pattern entity*). Pattern entities of a pattern archetype exhibit a similar color scheme and overall shape but differ in scale, orientation, reflection and also non-affine transformations. For a surface painting to classify as a pattern archetype we require it to (i) have a finite extent on an object’s surface, excluding continuous friezes reaching around the whole solid of revolution, (ii) have at least two entities, and (iii) display some degree of ‘artistic richness’, omitting very simplistic paintings like delimiters or single-stroke patterns.

In terms of filing, we decided to save the information regarding the location and extent of a pattern entity as a set of face IDs of the triangle mesh the entity appears on. Hence, for pattern entities we devised the data structure

```
Pattern_Entity: {
  "face_ids": [uint32],
  "orientation": float32,
  "scale": float32,
  "reflection": bool
},
```

which additionally stores the orientation $o \in [0, 2\pi)$, scale $s \in (0, \infty)$ and reflection $r \in [0, 1]$ characteristics. The latter indicates whether or not the entity appears in a mirrored fashion with respect to the archetype’s characteristic example. Similarly, the data structure for the pattern archetype

```
Pattern_Archetype: {
  "object_id": uint32,
  "fold_symmetry": int32,
  "sample_entity": Image,
  "entities": [Pattern_Entity]
}
```

includes the information to which object the annotation belongs to, as well as an image of one representative entity and a list of associated pattern entities. The integer-valued property `fold_symmetry` stores the archetype’s rotational symmetry order n , since the shapes of many archetypes have one or more symmetry axes (e.g., the sun wheel in Fig. 1 left). For those which do not exhibit any symmetry, this property is set to 0. Annotations are stored on the file system using the JavaScript Object Notation (JSON) format. This has two

major advantages. Firstly, the devised data structure can be readily adjusted and extended to incorporate additional annotation properties. Secondly, the JSON format is well supported with parser for all established programming languages, allowing fellow researchers to process the annotations with their own tools if they prefer to do so.

Our first experiments unveiled that the mesh resolution obtained from the structured-light scanner was not high enough to support a precise per-face surface labeling. Hence, the raw meshes have been subjected to one refinement step of Loop’s subdivision [Loo87], resulting in about 130K faces per 3D model. Other preprocessing steps include the removal of inner surface areas (if present), the removal of non-manifold edges together with their faces, and the removal of unreferenced vertices.

5. Annotation Tools

The process of generating annotations for 3D models is a primarily manual task requiring expert knowledge, which we support with custom made tools. The only required input is the textured 3D mesh for which the annotation is created and the resulting output is a JSON file containing the annotation data structure given in Sec. 4. We decided to divide process over two inter-operating tools which has the advantage that the annotation effort can be easily split up between different groups of people. The first tool (*Sample Pattern Selection*, Sec. 5.1) is responsible for determining all pattern archetypes present on the input object and outputs a set of images, each displaying a characteristic entity of those while the second tool (*Pattern Entity Selection*, Sec. 5.2) takes these images and the 3D mesh as input and yields the annotation file (Fig. 4).




Figure 4: Annotation pipeline with input (3D mesh), output (annotations in JSON format) and intermediate results.

5.1. Sample Pattern Selection

Let \mathcal{S} be the surface of a 3D object. To specify an archetype pattern, we ask the user to select a set of surface points $\mathbf{P} \in \mathcal{S}$ around one of its entities. The first step is fitting a plane $T = \mathbf{n} \cdot \mathbf{x} + d$ to the set \mathbf{P} , and determining its outward-pointing normal \mathbf{n} (Fig. 5, left). We then project the centroid \hat{P} of \mathbf{P} onto \mathcal{S} along the normal direction \mathbf{n} to obtain the central selection point \hat{S} on the surface. To this end, the ray $\hat{P} + \mathbf{n}t$ is intersected with \mathcal{S} . The largest intersection parameter t_{max} then defines the surface point $\hat{S} = \hat{P} + \mathbf{n}t_{max}$.

In the second step, we set up the camera to obtain an encompassing image of the selected pattern entity that is centralized in \hat{S} (Fig. 5, right). To this end, we create a perspective camera with 90° vertical field of view and 256×256 resolution. To obtain the suitable view transformation, the shape is first translated to place

\hat{S} in the 3D origin and rotated such that the normal vector \mathbf{n} aligns with the Z+ axis. The camera origin is then offset along \mathbf{n} by a distance $c \cdot m$, where $m = \max_{P \in \mathbf{P}} \|P_{xy} - \hat{P}_{xy}\|$ denotes the maximum distance of points from the view axis and $c = 1.2$ considers some padding around the pattern in the image.




Figure 5: Left: A point \hat{S} on the surface S is computed as the central surface point of the selection. Right: Camera setup for capturing the corresponding pattern selection image.

5.2. Pattern Entity Selection

For the selection of all entities of a given pattern archetype we decided to circumnavigate the complex task of surface annotation in 3D by projecting the curved object surface to a planar proxy. Due to the rotational symmetry exhibited by all of the 3D models we applied a variation of the cylindrical unwrapping by Karras et al. [KPP96] to that end. As the model's surface needs to be 'cut open' at one point along the angular direction, pattern entities in the vicinity of this cut might be severed. Hence, we added additional padding to both ends of the planar surface, by copying the surface, to allow for uninterrupted annotation of all pattern entities. Due to the jagged appearance of most patterns a polygon lasso tool proved to be the most convenient tool for selecting a set of faces. All faces with at least one vertex inside the polygon are determined efficiently with a flood-fill approach.

The GUI we devised for the pattern entity selection is tripartite and includes (i) a color-coded selection list of the loaded pattern archetypes with a thumbnails showing the sample images (Fig. 6①), (ii) a navigable graphics widget displaying the unwrapped object surface (Fig. 6②), and (iii) a 3D viewer, previewing the current annotation on the 3D model (Fig. 6③). When hovering over the unwrapped surface, the 3D model aligns such that the hovered-over position is facing the camera.

The annotation process is conducted by loading a 3D mesh, which is subsequently unwrapped, as well as a set of pattern archetypes belonging to this object, given by the sample images obtained from the *Sample Pattern Selection*. The following steps are conducted iteratively by the user until all pattern entities of all pattern archetypes are annotated:

1. (Optional) Select a pattern archetype from the selection list (Fig. 6①) if the pattern entity you want to select next does not belong to the currently selected pattern archetype.

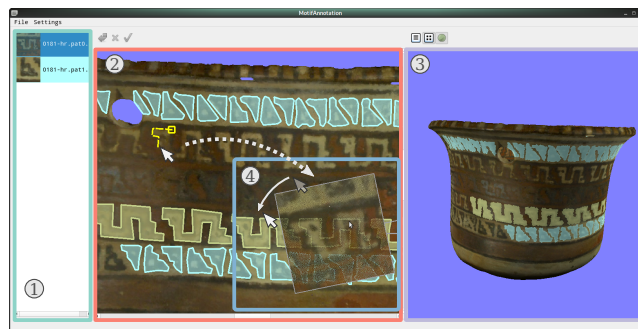


Figure 6: Sample Pattern Selection with (1) Colored selection list, (2) Occurrence selection on flattened surface, (3) 3D viewer with annotated model.

2. Select a pattern entity on the flat proxy surface with a polygonal lasso tool (Fig. 6②).
3. The Pattern Entity Selection switches into entity edit mode which allows to specify the orientation and scale with respect to the archetype's sample image by aligning it with the current selection (Fig. 6④). At this point, it can also be declared whether or not the current entity is a reflection with respect to the archetype sample.

At the end of the annotation process the result can be exported as a JSON file. The tool also allows to load and edit existing annotations. While the appropriate selection of sample patterns (Sec. 5.2) is oftentimes debatable and requires domain knowledge, the selection of the individual entities is usually straightforward and can be conducted by laymen.

6. Evaluation Metrics

We evaluate the benchmark on a per-pattern archetype level (Sec. 4) and assume that a retrieval algorithm is provided with a query, constituting either the archetype's sample image or one of its pattern entities. Given this input, the retrieval is required to return a set of coherent surface areas $\mathbf{R} = \{\mathcal{R}_j\}, j \in J$ (detections), each of which given by a set of face ids \mathcal{R}_j , and referring to a detected pattern entity. J denotes the index set of retrieval results. Alongside the detections some retrieval algorithms also yield associated confidence values. The evaluation metric we propose (Sec. 6.1) follows the method presented in the PASCAL Visual Object Classes Challenge [EEVG*15] for object detection tasks.

For techniques that approach the retrieval problem in 2D space after a surface mapping (c.f. Pattern Entity Selection in Sec. 5.2), we provide an alternative evaluation metric that allows to evaluate the effectiveness of a retrieval method in image space. To this end, we assume that such a retrieval algorithm yields 2D point locations and – optionally – rotation, scale and reflection transformation with respect to the query (Sec. 6.2).

For both metrics implementations are provided on the benchmark website which can be readily employed to evaluate the performance of different retrieval algorithms. Note that the presented metrics are mere suggestions and that researchers are invited to develop own metrics to evaluate other traits of retrieval methods.

6.1. Mean Average Precision (MAP)

Let $\mathbf{G} = \{\mathcal{G}_i\}, i \in I$, denote the set of surface areas corresponding to the pattern entities of the given archetype (ground truth). We determine a correspondence between the elements of \mathbf{R} and \mathbf{G} and compute the Intersection-over-Union (IoU) rate for every possible unordered pair combination $\{\mathcal{R}_j, \mathcal{G}_i\}$. Since there could be detections that have a considerable overlap with more than one ground truth annotation, we need to assign the detections to avoid multiple correspondences. We proceed greedily by first sorting the potential pairings according to their decreasing IoU rates. Thus, our method iterates over the assigned pairs while selecting one pair at a time, and the detections and patterns involved in the selection cannot be chosen again in a subsequent pair. This procedure yields a correspondence set between detections and pattern occurrences.

Next, our method sorts the correspondence set by the confidence of the detections. We also mark the pairs with IoU above a threshold t_{MAP} as true positives, and detections with IoU below t_{MAP} as false positives. The ground truths without a match in the detections become false negatives. We use the ranking list of correspondences to compute recall and precision values, and subsequently calculate the 11-points interpolated precision values for the archetype pattern. Finally, the average precision for an archetype is the average of the eleven interpolated precision values. The final evaluation measure – MAP – is the mean value over the average precisions of patterns in the complete collection. In all the experiments in our paper, the threshold for IoU is $t_{MAP} = 0.5$.

6.2. Average Nearest-Neighbor Distance (ANND)

The ANND metric operates on the flattened object surface and returns a normalized average distance. The method can be applied for methods which return the center point locations of detections $\mathbf{C} = \{C_j\}_{j \in J} \subset \mathbb{R}^2$ and optionally also the affine transformations scale $S = \{s_j\} \subset \mathbb{R}_+$, orientation $O = \{o_j\} \subset \mathbb{R}_+$ and reflection $R = \{r_j\} \subset \mathbb{B}$, with respect to a given query. Hence, $\{C_i\}_{i \in I}, \{s_i\}, \{o_i\}$ and $\{r_i\}$ describe the location, scale, orientation, and reflection of the ground truths respectively. Our proposed ANND metric is given by

$$\bar{d}_{ANND} = \begin{cases} \frac{1}{|J|} \sum_{i=1}^{|I|} \bar{d}(i, \hat{\phi}(i)) + (|J| - |I|), & \text{if } |I| \leq |J| \\ \frac{1}{|I|} \sum_{j=1}^{|J|} \bar{d}(\hat{\phi}^{-1}(j), j) + (|I| - |J|), & \text{otherwise} \end{cases} \quad (1)$$

with $\bar{d}(i, j)$ as the normalized distance between the ground truth i and the detection j , which is given by $\bar{d} = \tilde{d}_G$, for retrieval algorithms which return exclusively point locations \mathbf{C} . For retrievals algorithms which also yield orientation, scale, and reflection information, $\bar{d} = (\tilde{d}_G + \tilde{\Delta}_S + \tilde{\Delta}_O + \tilde{\Delta}_R) / 4$ may be used. Note that the accumulated distances of the optimal assignment $\hat{\phi} = \text{argmin}_{\phi} \sum_{i=1}^{|I|} d_{i\phi(i)}$ of detections to ground truths is used, to prevent that multiple detections are mapped to one and the same ground truth. While $\bar{d} \in [0, 1]$ yields the normalized distance between matched detections and ground truths, the second term $|J| - |I|$ or $|I| - |J|$ in Eqn. (1) serves as a penalty if either too many or not sufficient detections are made.

The geometric distance $\tilde{d}_G(i, j) = 1 - \theta(\|C_i - C_j\|, h_i)$ with

the fast decaying smooth function $\theta(r, h) = e^{-r^2/(h/2)^2}$ ascertains that the distance descends towards zero for a good detection and is close to one otherwise. The respective support radii $\{h_i\}_{i \in I} \subset \mathbb{R}_+$ are given by the circumradii of annotation polygons. $\tilde{\Delta}_S(i, j) = \max\{|s_i - s_j|, \Delta_{S_{max}}\} / \Delta_{S_{max}}$ describes the similarity of scales which is capped with $\Delta_{S_{max}} = 2$. The similarity of orientations is given by

$$\tilde{\Delta}_O(i, j) = \begin{cases} \min_{\psi \in \Psi} \{d_{\angle}(o_i + \psi, o_j)\} / \pi, & \text{if } n > 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

with n as the n -fold rotational symmetry of the pattern archetype, d_{\angle} as the absolute difference between two angles and $\Psi = \{2\pi k/n : k \in [0..n-1]\}$ as the orientations of the archetype's n symmetry axes. The similarity of reflection states is defined as

$$\tilde{\Delta}_R(i, j) = \begin{cases} 0, & \text{if } r_i = r_j \\ 1, & \text{otherwise.} \end{cases} \quad (3)$$

7. Experimental Evaluation

Even though there are no retrieval algorithms tailored to this specific problem yet we evaluate the dataset with (i) purely random-generated results (Sec. 7.1) to determine a lower baseline, (ii) a very naïve approach leveraging bitwise surface comparison at randomized locations (Sec. 7.2), and (iii) a domain related algorithm (Sec. 7.3) which is designed for binary patterns but still applicable for some surfaces. Note that these approaches operate on the flat surface image, but in order to evaluate the MAP metric (Sec. 6.1) the predicted locations and surface areas have been remapped to the 3D mesh. For each pattern archetype one retrieval, based on a randomly selected entry $q \in I$ as query, is conducted.

7.1. Random Selection (RND)

As the name suggests, with this approach we generate pattern detections with individual locations, as well as orientation scale and flip-state, purely at random. Per retrieval $N_p \sim \text{Poisson}(\mu_{entities})$ detections are generated, with $\mu_{entities} = 24.8$ as the mean number of pattern entities per archetype. The detections' locations $\mathbf{C} \sim \mathcal{U}((0, 0)^T, (w, h)^T)$ (with w as the flat surface width and h as the flat surface height), orientations $O \sim \mathcal{U}(0, 2\pi)$, and flip-states $R \sim \mathcal{U}\{0, 1\}$ follow uniform distributions, while their scales $S \sim \mathcal{N}(1, \sigma_{scale})$ follow a normal distribution with $\sigma_{scale} = 0.08$ empirically determined from the dataset.

7.2. Randomized Local XOR (RLXOR)

RLXOR is a slightly improved version of RND. For the given query $N_0 = 4K$ initial detections are generated with the same distributions as in Sec. 7.1. For each detection $j \in J$ a confidence value

$$\kappa_j = \frac{1}{|S_q|} \sum_{k=1}^{|S_q|} S_q[k] \oplus S_j[k] \quad (4)$$

is computed with S_q as the set of ordered RGB surface pixels which are within a radius of $r = s_j r_{query}$ from the location of the query C_q and S_j as the equal-sized set of surface pixels around the location of the detection C_j . The detections are sorted by their confidence values $\{\kappa_j\}$, before the top- $N_p \sim \text{Poisson}(\mu_{entities})$ are returned.

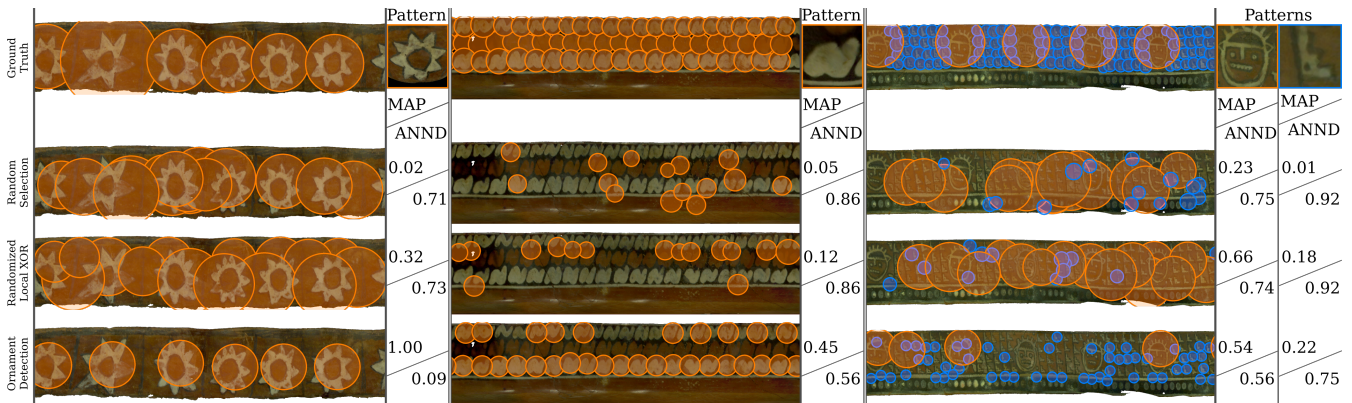


Figure 7: The comparison of different retrieval results and evaluation metrics with three objects with varying complexity of surface patterns.

7.3. Repetitive Ornament Detection (ROD)

Lengauer et al. [LKK⁺20] propose a recognition method for repetitive patterns which computes a confidence map of the dimension of the flattened object by XOR-ing a binarized query depiction with the binarized flat object surface. This confidence map is further refined by leveraging certain traits, intrinsic to ancient painted pottery. Even though the approach is tailored to Greek pottery which exhibits bichrome surface paintings, it has also shown promising results with polychromatic surface paintings if they can be binarized such that patterns are preserved.

7.4. Application on the dataset

We conducted a qualitative evaluation of our metrics with three different objects from our dataset, exhibiting different quantities of patterns and different degrees of complexity. The first object displays six entities of a star-shaped pattern (Fig. 7, left), with slightly varying numbers of jags and overall shape. As expected the retrieval with RND performs very poorly according to both metrics. The RLXOR performs equally poor, at least according to the ANND, which was found to be due to the too large number of detections. With ROD an almost perfect score was achieved with both metrics which was attributed to the relative simplicity of the object's surface painting and the binary colorization. The second object (Fig. 7, middle) displays 65 very similar 'N'-shaped patterns arranged in grid structure. As with the previous example, RND and RLXOR perform rather poorly as the number of detections differs greatly from the number of entities in the ground truth. The results obtained with ROD is approximately at the center of the scores of both metrics ($[1, 0]$ for ANND and $[0, 1]$ for MAP), which meets our expectations as half of the bicolored surface patterns was lost during the binarization necessary for this approach. The third object (Fig. 7, right) poses the most complex challenge and exhibits two distinct pattern archetypes: a rather large face shape with five entities and smaller triangle shape with 113 entities. With this example a slight improvement of RLXOR over RND is observable. With ROD half of the entities of the first pattern are correctly detected and about a third of the second, which is clearly reflected in the evaluation scores.

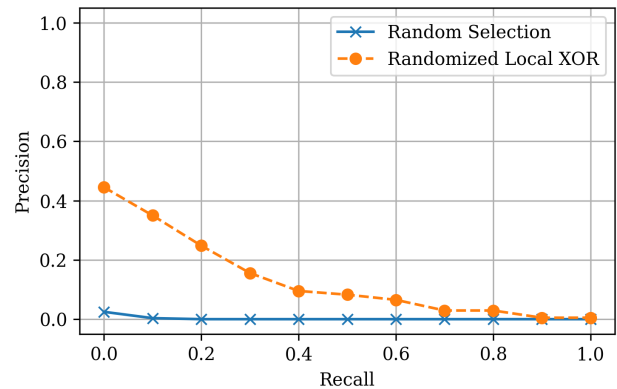


Figure 8: The precision vs. recall curves for RND and RLXOR.

Additionally, we performed a quantitative evaluation on the whole dataset. Fig. 8 shows the precision vs. recall curves for the implemented pattern detection algorithms. Generally, these results indicate that there is huge room for improvements. As expected, the trivial approach (RND) yields a very low precision for all recall levels. The second approach (RLXOR) proves to be more effective than Random selection, but is still far from optimal. Indeed, even for a recall of 0 the obtained precision is lower than 0.5 with this method, which is quite low for practical applications. The third approach (ROD) was not able to yield results for many of the objects as it requires a binarized version of the surface texture. Since this is not feasible for most of the polychrome paintings, this approach is not included in the quantitative equation. These results show that finding repetitive patterns on ancient pottery surfaces is non-trivial and requires the development of novel and robust solutions for achieving a high effectiveness on this task.

8. Possible Extensions

The current selection of objects in the datasets solely contains ancient Peruvian pottery. However, the annotation tools provided along with the dataset (Sec. 5) can be easily used to extend the

collection with artifacts from other cultures, e.g., ancient Greek pottery, or any kind of textured 3D meshes exhibiting repetitive surface patterns. The flexibility of our data structure and the JSON file format used for storing allows to add additional properties, relevant for other models. Even though patterns on Peruvian pottery generally do not exhibit any identifiable hierarchical or symmetry-based relations, those could be added on top of our low-level annotation by introducing pair-wise relations between pattern entities or archetypes, such as ‘is twin of’, ‘has’, or ‘belongs to’. An example is given by the three-pronged cactus shape shown in the inset, which appears several times with three blossom shapes.



Extensions are not limited to extensions of the dataset and their annotations. Note that the evaluation metrics we propose are just a starting point. They do not capture any efficiency properties of a method, for instance. Hence, arbitrary additional metrics, tailored to specific requirements of a retrieval system can be defined.

9. Conclusion

We present an annotated dataset of textured meshes with repetitive surface patterns which can be readily extended with our custom made annotation tools. The automatic recognition of these surface patterns poses a non-trivial problem which can be tackled with geometric approaches based on the textured 3D mesh or with image-based methods operating on a 2D parametrization of the surface. To this end, we propose two orthogonal evaluation metrics and demonstrate their applicability by evaluating a naïve sampling-based pattern detection technique as well as a domain related approach.

Acknowledgments

This work was co-funded by the Austrian Science Fund FWF and the State of Styria, Austria within the project *Crossmodal Search and Visual Exploration of 3D Cultural Heritage Objects* (P31317-NBL). This work has been also partially supported by Proyecto de Mejoramiento y Ampliación de los Servicios del Sistema Nacional de Ciencia, Tecnología e Innovación Tecnológica (Banco Mundial, Concytec), Nr. Grant 062-2018-FONDECYT-BM-IADT-AV. This work was also co-funded by ANID - Millennium Science Initiative Program - Code ICN17_002. Special thanks to Diana Vargas and Norma Menacho for conducting the scanning in the Josefina Ramos de Cox museum, in Lima, Perú.

References

- [BM17] BOGACZ B., MARA H.: Automatable annotations—image processing and machine learning for script in 3d and 2d with gigamesh. *Kodikologie und Paliographie im Digitalen Zeitalter* (2017), 137–149. 3
- [BMTB*18] BIASOTTI S., MOSCOSO THOMPSON E., BARTHE L., BERRETTI S., GIACHETTI A., LEJEMBLE T., MELLADO N., MOUSTAKAS K., MANOLAS I., DIMOU D., TORTORICI C., VELASCO-FORERO S., WERGH N., POLIG M., SORRENTINO G., HERMON S.: Recognition of Geometric Patterns Over 3D Models. *Eurographics Workshop on 3D Object Retrieval* (2018), 7 pages. Artwork Size: 7 pages ISBN: 9783038680536 Publisher: The Eurographics Association. 2
- [BWM*11] BERNER A., WAND M., MITRA N. J., MEWES D., SEIDEL H.-P.: Shape analysis with subspace symmetries. *Computer Graphics Forum* 30, 2 (2011), 277–286. 3
- [EEVG*15] EVERINGHAM M., ESLAMI S. M. A., VAN GOOL L., WILLIAMS C. K. I., WINN J., ZISSERMAN A.: The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision* 111, 1 (Jan. 2015), 98–136. 5
- [GCO06] GAL R., COHEN-OR D.: Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph.* 25, 1 (Jan. 2006), 130–150. 3
- [HGM14] HUANG Q., GUIBAS L. J., MITRA N. J.: Near-regular structure discovery using linear programming. *ACM Trans. Graph.* 33, 3 (June 2014). 3
- [IT11] ITSKOVICH A., TAL A.: Surface partial matching and application to archaeology. *Computers & Graphics* 35, 2 (2011), 334–341. Virtual Reality in Brazil Visual Computing in Biology and Medicine Semantic 3D media and content Cultural Heritage. 3
- [KBW*12] KALOJANOV J., BOKELOH M., WAND M., GUIBAS L., SEIDEL H.-P., SLUSALLEK P.: Microtiles: Extracting building blocks from correspondences. *Computer Graphics Forum* 31, 5 (2012), 1597–1606. 3
- [KPP96] KARRAS G., PATIAS P., PETSAS E.: Digital monoplotting and photo-unwrapping of developable surfaces in architectural photogrammetry. *International Archives of photogrammetry and Remote Sensing* 31 (1996), 290–294. 5
- [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), Eurographics Assoc. 3, 7
- [Loo87] LOOP C.: Smooth subdivision surfaces based on triangles. *Master’s thesis, University of Utah, Department of Mathematics* (1987). 4
- [MBB10] MITRA N. J., BRONSTEIN A., BRONSTEIN M.: Intrinsic regularity detection in 3d geometry. In *Computer Vision – ECCV 2010* (Berlin, Heidelberg, 2010), Daniilidis K., Maragos P., Paragios N., (Eds.), Springer Berlin Heidelberg, pp. 398–410. 3
- [MTW*18] MOSCOSO THOMPSON E., TORTORICI C., WERGH N., BERRETTI S., VELASCO-FORERO S., BIASOTTI S.: Retrieval of Gray Patterns Depicted on 3D Models. *Eurographics Workshop on 3D Object Retrieval* (2018), 7 pages. Artwork Size: 7 pages ISBN: 9783038680536 Publisher: The Eurographics Association. 2, 3
- [TB18] THOMPSON E. M., BIASOTTI S.: Edge-based LBP Description of Surfaces with Colorimetric Patterns. *Eurographics Workshop on 3D Object Retrieval* (2018), 8 pages. Artwork Size: 8 pages ISBN: 9783038680536 Publisher: The Eurographics Association. 3
- [TBDC19] THOMPSON E. M., BIASOTTI S., DIGNE J., CHAINE R.: mpLBP: An Extension of the Local Binary Pattern to Surfaces based on an Efficient Coding of the Point Neighbours. *Eurographics Workshop on 3D Object Retrieval* (2019), 8 pages. Artwork Size: 8 pages ISBN: 9783038680772 Publisher: The Eurographics Association Version Number: 009-016. 3
- [TBG*20] THOMPSON E. M., BIASOTTI S., GIACHETTI A., TORTORICI C., WERGH N., OBEID A. S., BERRETTI S., NGUYEN-DINH H.-P., LE M.-Q., NGUYEN H.-D., TRAN M.-T., GIGLI L., VELASCO-FORERO S., MARCOTEGUI B., SIPIRAN I., BUSTOS B., ROMANELIS I., FOTIS V., ARVANITIS G., MOUSTAKAS K., OTU E., ZWIGGELAAR R., HUNTER D., LIU Y., ARTEAGA Y., LUXMAN R.: SHREC 2020: Retrieval of digital surfaces with similar geometric reliefs. *Computers & Graphics* 91 (2020), 199–218. 2
- [ÚSN*20] ÚBEDA I., SAAVEDRA J. M., NICOLAS S., PETITJEAN C., HEUTTE L.: Improving pattern spotting in historical documents using feature pyramid networks. *Pattern Recognition Letters* 131 (2020), 398–404. 3