Stefan Lengauer*, Peter Houska, Reinhold Preiner, Elisabeth Trinkl, Stephan Karl, Ivan Sipiran, Benjamin Bustos, and Tobias Schreck

# Interactive annotation of geometric ornamentation on painted pottery assisted by deep learning

https://doi.org/10.1515/itit-2022-0007
Received January 23, 2022; revised July 9, 2022; accepted July 21, 2022

**Abstract:** In Greek art, the phase from 900 to 700 BCE is referred to as the Geometric period due to the characteristically simple geometry-like ornamentations appearing on painted pottery surfaces during this era. Distinctive geometric patterns are typical for specific periods, regions, workshops as well as painters and are an important cue for archaeological tasks, such as dating and attribution. To date, these analyses are mostly conducted with the support of information technology. The primitives of an artefact's ornamentation can be generally classified into a set of distinguishable pattern classes, which also appear in a similar fashion on other objects. Although a taxonomy of known pattern classes is given in subject-specific publications, the automatic detection and classification of surface patterns from object depictions poses a non-trivial challenge. Our long-term goal is to provide this classification functionality using a specifically designed and trained neural network. This, however, requires a large amount of labelled training data, which at this point does not exist for this domain context. In this work, we propose an effective annotation system, which allows a domain expert to interactively segment and label parts of digitized vessel surfaces. These user inputs are constantly fed back to a Convolutional Neural Network (CNN), enabling the prediction of pattern classes for a given surface area with ever increasing precision. Our work paves the way for a fully automatic classification and analysis of large surface pattern collections, which, with the help of suitable visual analysis techniques, can answer research questions like pattern variability or change over time. While the capability of our proposed annotation pipeline is demonstrated at the example of two characteristic Greek pottery artefacts from the Geometric period, the proposed methods can be readily adopted for the patternation in any other chronological periods as well as for stamped motifs.

**Keywords:** Painted pottery, ornament detection and classification, ornamentation benchmark, evaluation, visual clustering

**ACM CCS:** Applied computing → Arts and humanities, Computing methodologies → Artificial intelligence → Computer vision → Computer vision problems → Matching, Information systems → Information retrieval → Specialized information retrieval → Multimedia and multimodal retrieval → Image search, Human-centered computing → Visualization → Visualization application domains → Visual analytics

***Corresponding author: Stefan Lengauer,** Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Inffeldgasse 16c, 8010 Graz, Austria, e-mail: s.lengauer@cgv.tugraz.at
**Peter Houska, Reinhold Preiner, Tobias Schreck,** Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Inffeldgasse 16c, 8010 Graz, Austria, e-mails: p.houska@cgv.tugraz.at, r.preiner@cgv.tugraz.at, tobias.schreck@cgv.tugraz.at
**Elisabeth Trinkl, Stephan Karl,** University of Graz, Institute of Classics, Universitätsplatz 3, 8010 Graz, Austria, e-mails: elisabeth.trinkl@uni-graz.at, stephan.karl@uni-graz.at
**Ivan Sipiran, Benjamin Bustos,** University of Chile, Department of Computer Science, Beauchef 851 3er piso, Santiago, Chile, e-mails: isipiran@dcc.uchile.cl, bebustos@dcc.uchile.cl

## 1 Introduction

The study of ancient Greek pottery is an essential research topic in the field of classical archaeology. Pottery, preserved in a huge number due to the relative durability of the ceramic material, is a key to answer questions regarding society, economy, daily life and religion of ancient populations. Archaeological research on these objects involves not only the morphological classification of the object itself, but also the analysis of their surface paintings, since most of the ancient Greek pottery is elaborately painted. Long-lasting research on these decorations, so-called vase paintings, have revealed different techniques and styles indicating chronological periods, workshops or painters. One specific group of Greek pottery is referred to as *Geometric* due to their characteristic surface paintings which resemble an assemblage of various basic geometric shapes like straight lines, circles and rectangles. Geometric pottery belongs to the second period of the Greek Dark Ages, the *Geometric period*, lasting from ca. 900 to 700 BCE. Most of the Geometric pottery is densely covered by linear ornament patterns. A categorization of these patterns helps to discover similarities and relations between objects from different sites or collections and to reveal, for example, sets of objects originating from the same region, workshop or painter. The identification of such object relations allows domain experts to gain new insights, e. g., concerning the distribution or workshop identifications.

To date, object collections are mostly available in the form of domain publications where the information for a specific object is available as a set of photographs together with textual descriptions and – in rare cases – archaeological drawings of the surface paintings. The descriptions however are generally non-standardized and unstructured, describing either the object as a whole or referencing certain parts, e. g., a vessel's shoulder, but lacking any direct reference to specific image regions. References to similar patterns on other objects are only given in the textual description with predicates such as 'similar to X' or 'as in Y', lacking the possibility to reconstruct which aspects were responsible for the stated similarity. Recent publications may also contain references to Kunisch [1], who presents a very comprehensive, yet not all encompassing, taxonomy of known pattern classes present on Geometric pottery objects.

However, in course of the study of the specific decoration on pottery to determine when and where it was made, archaeologists have to conduct numerous pairwise comparisons between an object in question and plausible candidates, optimally from dated find contexts like graves or closed deposits. To qualify pattern as similar, domain experts take into account not only the pattern's overall appearance, but countless other minuscule details, e. g., alternating stroke directions in the hatching of a meander. When analyzing an object's surface paintings, experts typically determine: (i) which distinguishable pattern classes are present, (ii) which given patterns are similar to those described by Kunisch [1], (iii) which pattern classes appear together on the surface and (iv) on which part of the vessel the patterns are located, e. g., neck, body, handle, etc. Fig. 1 shows a thorough analysis for the flattened surface of the Attic Geometric pitcher *Vienna KHM IV 1*. An apparent disadvantage of this purely manual workflow is the entailing time investment, resulting in many objects ending up without proper surface documentation, leading to possibly countless undetected object relations. Moreover, without structured annotation and labelling, searching for

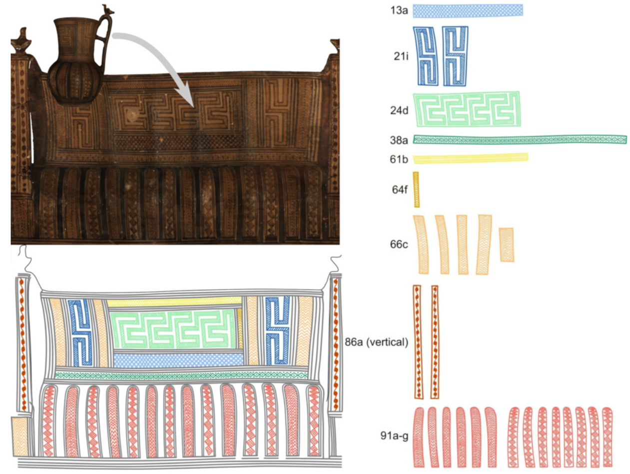


**Figure 1:** Painted object surface after unwrapping the scanned 3D model (top left), in-place graphical annotation of the flat surface by archaeologist (bottom left) and the grouping and classification of individual patterns (right).

combinations of patterns on surfaces is very complicated or not even possible at all.

We propose a structured, computer-aided surface annotation pipeline which requires minimal effort from the user and ensures quantitative reproducibility of annotations through well designed similarity metrics. To this end, a form of automatic detection is required.

A group of baseline computer vision methods for determining descriptive numerical representations from image data are referred to as *engineered features*. Those methods aim to formulate numerical representations of certain traits of an image. Even though they have been successfully applied for a variety of tasks over decades it is often hard to formulate – and anticipate – the minuscule details required for a reasonable detection and classification and oftentimes suffer from robustness issues in the face of eroded and incomplete patterns. In contrast, *learned features* rely on algorithms which improve automatically through the supply of respective training data. Convolutional Neural networks (CNN) in particular have received ever increasing attention the last years and constitute the state-of-the art for image detection and classification tasks [2]. However, they generally require a large amount of labelled training data to provide meaningful results, which currently do not exist in sufficient quantity for the domain problem relating to Greek Geometric pottery.

Hence, our idea is that: **(1)** an annotation pipeline in its initial state resembles the current manual analysis process, and **(2)** a CNN in the background learns from every action the user performs, gradually improves, and **(3)** eventually converges to an expert system which can recognize pattern elements at high precision.

The remainder of the paper is structured as follows: Sec. 2 gives an overview on related work on texture recognition on CH objects, interactive labeling and related benchmark datasets; Sec. 3 discusses our proposed annotation pipeline in depth; Sec. 4 shows first results which are discussed in Sec. 5; and finally Sec. 6 discusses further potential applications of an annotated dataset.

## 2 Related work

Our proposed annotation pipeline constitutes an interactive online learning approach. Hence, we first of all discuss related online learning approaches, before we have a look at pattern recognition techniques, since our system should be able to propose a fitting pattern class for a user selection. Finally, we show several established benchmarks for textual and geometric surface patterns.

#### Interactive machine learning and labeling

Involving the 'human in the loop' of the machine learning process is a valuable goal for many reasons, e. g., understanding of data and learned rules by human decision makers, providing training data, or interactively adjusting the parameters of the machine learning models. Also, explainability and trust in machine learning methods are increasingly important, as more and more decisions are nowadays automatized [3]. To date, many approaches to integrate machine learning and interactive approaches have been proposed. Visual analytics approaches rely on visualization to show data, models, parameters and allow for interaction [4]. Relevant to our work are previous approaches in interactive labeling [5, 6], which aim to provide a small number of training data items which improve the machine learning models, i. e., classification tasks in these works. Typically, a number of views allow the user to inspect data features and classification boundaries of the trained methods. Users can explore that data, change or add labels, and compare the effects of their labeling on the machine learning models. While the considered data can be abstract and high-dimensional if suitable visualization techniques are used [5, 6], the elements to be labeled are image-based, i. e., patterns on 3D objects, in our proposed pipeline.

#### Pattern recognition

Lengauer et al. [7] present a semi-automated annotation of repetitive ornaments on 3D pottery surfaces. They leverage the observation that filling ornaments often appear in a repetitive fashion around the solid of revolution. To this end, they combine a naïve sliding window approach with a probability weighting based on the number of detections at a certain point along the axis of rotation. While this approach is applicable for simple ornaments which appear repetitively, it fails for the segmentation and classification of Geometric pottery where individual patterns oftentimes appear just a single time on an object surface.

Zhou et al. [8] present a web-based system for matching pottery sherds, exhibiting stamped reliefs on their surface, to the respective paddle stamp design that was most likely originally used for making these specific imprints. For each 3D scanned sherd uploaded by a user, a CNN extracts a binary image that encodes the (partial) curve patterns of the stamp design. Another CNN consumes these binary images and determines the best matching reference stamp design stored in a remote database. Curve-pattern extraction and sherd-to-design matching are performed remotely on a GPU cluster. This allows archaeologists to use the service directly from their browser for remote collaboration.

Crowley and Zisserman [9] describe a system that is able to automatically annotate depictions of gods and animals on vases. They processed a large pottery database that not only features several photographs for each documented vase, but also stores brief descriptions of the depicted scenes. By applying text-mining methods they are able to link the featured gods and animals mentioned in these descriptions to the actual image area in which they appear. Using this knowledge, they label new images with appropriate descriptions of which figures can be seen on the painted pottery surface.

#### Benchmark datasets

Several entries have been made to the 3D Shape Retrieval Challenge (SHREC[1]). The SHREC 2017 track "Retrieval of surfaces with similar relief patterns" by Biasotti et al. [10] proposes a benchmark for retrieving surface patches from a 3D model with a relief pattern similar to a given example. The dataset consists of 15 classes of the so-called "tissues", obtained from different fabrics, which define each individual relief pattern of interest. Those relief patterns are applied to non-textured 3D models (12 models per class), producing a dataset of 180 textured models. I. e., the relief patterns are obtained from real-world fabrics, while the 3D models are synthetically generated. The experimental evaluation shows that the overall retrieval precision on

1 http://www.shrec.net/

this benchmarks is rather low, indicating the task is a difficult challenge.

The SHREC 2018 track "Recognition of geometric patterns over 3D models", also by Biasotti et al. [11], proposes a benchmark for recognizing relief patterns on 3D models. The dataset is built from laser-scanned fragments of archaeological artifacts, comprising 38 triangle meshes divided in "query dataset" (8 meshes) and "model set" (30 meshes). The benchmark proved to be particularly difficult as the organizers of this SHREC track received no satisfactory result by any algorithm or method. The main difficulties reported by the groups participating on this track were: dealing with noisy, real-world data; dealing with high-definition data; and lack of training datasets.

The SHREC 2018 track "Retrieval of gray patterns depicted on 3D models" by Moscoso Thompson et al. [12] proposes a benchmark for retrieving objects with a given painted texture pattern. For this benchmark, the dataset is synthetically generated using binary patterns. The 3D models can be fully covered by a single pattern, can have a faded double pattern, or can be partially covered with the pattern. The organizers define two different datasets: Single Pattern (200 models) and Complete Dataset (100 additional models with double textures). The goal of the task is to group the objects according to their texture pattern, while disregarding the objects' the geometry.

The SHREC 2020 track "Retrieval of digital surfaces with similar geometric reliefs" by Moscoso Thompson et al. [13] proposes a benchmark for retrieving surface patches with similar geometric relief. The global geometry of 3D models is not relevant for this task. The dataset is synthetically generated with 11 different geometric reliefs, introduced to 20 base 3D models. The patterns are purely geometric and are applied to the whole surface of the 3D model. The goal of this track is to group the models w. r. t. the different geometric reliefs. The experimental evaluation shows that the best methods for this task are based on neural networks. However, some methods based on engineered descriptors also obtained a high effectiveness score.

Lengauer et al. [14] present a dataset and a benchmark for repetitive pattern recognition on textured 3D surfaces. The dataset consists of 82 scanned painted ancient Peruvian vessels from the Josefina Ramos de Cox Museum in Lima, Peru. The main characteristics of the dataset are: it contains polychromatic texture patterns; the artifacts exhibit different levels of repetitiveness in their surface patterns; and, the patterns and motifs in the dataset can appear anywhere on the surface of the 3D model. All 3D models of this dataset are manually annotated by archaeologists, who identify clearly distinguishable pattern classes. To this end, they also manually annotate each occurrence of a repetitive pattern. Alongside the dataset itself, the authors also published the annotation tools[2] used for creating the dataset.

# 3 Interactive annotation

The design of our proposed pipeline assumes that users upload a digitized version (Sec. 3.2) of the object they want to annotate via a custom made annotation UI (Sec. 3.3). The annotation is conducted by alternating selection of a surface pattern, and assignment of the corresponding pattern class from a list. Upon selecting the surface area, the continuously growing database of pattern classes and annotations (Sec. 3.1) is scanned in the background, yielding a list of the most likely pattern classes, visible in the highlighted area, based on a similarity computation (Sec. 3.5) of features obtained from a CNN (Sec. 3.4). To support interoperability and easy collaboration among experts, we decide for a web-based application that connects to a common database. For straightforward integration of the annotations in other applications, we support the database export in a standardized structured manner.

## 3.1 Initial dataset

To support the traditional, manual expert annotation workflow, our system supports two distinct surface annotation tasks so that surface areas can be associated with both vessel parts and pattern classes.

### 3.1.1 Taxonomy of vessel parts

Our system offers a set of vessel part labels for referring to the different morphological parts of a pottery object. A taxonomy of these parts following the human body structure like foot, belly, neck is widely used in pottery studies [15]. We composed an initial label set $\mathcal{V}$, containing the most commonly used labels, beginning bottom-up: *foot*, *lower part*, *belly*, *shoulder*, *neck*, *rim* and *handle*. Yet, this set is not encompassing. Hence, we allow the user to dynamically add additional terms on demand.

2 https://datasets.cgv.tugraz.at/pattern-benchmark/

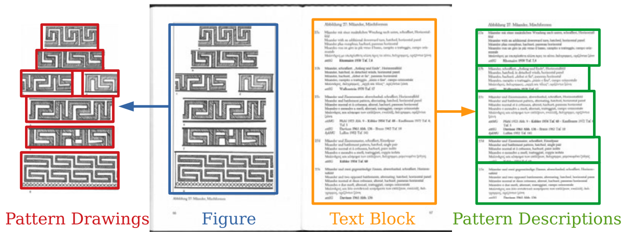


**Figure 2:** One characteristic page from the geometric pattern taxonomy by Kunisch [1]. The figures (blue) and text blocks (yellow) were detected and extracted automatically and further segmented into pattern drawings (red) and descriptions (green) respectively.

### 3.1.2 Taxonomy of geometric patterns

Our initial set of pattern classes is based on the taxonomy by Kunisch, presented in *Ornaments of Geometric Vessels* (in German) [1]. The volume comprises a total of 776 distinct pattern classes, collected from object depictions from more than 50 different domain publications. They have been grouped into 6 major – e. g. *rectangular ornaments* – and 20 minor groups. A specific pattern class is generally documented as:

- A representative black and white drawing of the pattern, usually closely resembling one of the related examples on a real artifact;
- A textual description, consisting of different labels referring to certain properties of the respective pattern, like 'hatched', 'vertical', etc., in up to five different languages; and
- (Optionally) literature references to publications depicting objects the patterns appears on.

This format is not very strict though, as several text descriptions refer to multiple pattern drawings.

Since the book was not available in digital format, the contents are extracted from the scans of a printed version with basic computer vision techniques. In a first step all text depictions are purged from the binarized pages with a morphological erosion, leaving only the figures containing the pattern drawings (Fig. 2, blue). After cutting out all figures and filling their respective rectangles with white, all remaining non-white areas were qualified as 'text blocks' (Fig. 2, yellow) and extracted as well. The figures were automatically segmented into the individual pattern drawings (Fig 2, red) by employing fundamental image processing techniques, implemented in the OpenCV computer vision library,[3] together with simple heuristics. The same procedure is used to segment the text blocks into individual pattern descriptions (Fig 2, green). The figure captions, combined with their respective subfigure captions, were leveraged to link the drawings with their respective descriptions. State-of-the-art Optical Character Recognition (OCR) methods [16] were are to transfer the rasterized text into machine readable form. All in all, about 90 % of entries were segmented, translated and linked correctly with these methods while the rest was fixed manually.

3 https://opencv.org

## 3.2 Input preparation

Our annotation system expects an image as input, which is the format in which the majority of CH objects are documented. Nonetheless, with advancements in scanning technology more and more objects are digitized as 3D models. Compared to images, which are limited to one view, those have the advantage that they are able to capture the complete and continuous object surface. To be processed by our pipeline the 3D model needs to be transferred to the 2D domain with appropriate methods. In the most simple case this could be a rendering, which however has the same downsides as a photograph. Hence, more sophisticated techniques create a planar projection of the 3D surface as a whole. This can be achieved by projecting the surface texture onto a shape primitive, e. g. cylinder, cone or sphere, fitted to the given geometry, which is then unwrapped virtually [17, 18]. The projective distortions introduced with this method can then be reduced with an elastic relaxation [19]. Houska et al. [20] demonstrated how this principle could also be applied to a set of overlapping photographs in order to reduce the distortions caused by the curved vessel surface.

## 3.3 Annotation interface

We design a web interface, where the uploaded image is displayed in the left half of the screen (Fig. 3 ①), while the right half of the screen shows a ranked list of the best matching pattern classes for the currently annotated area. Area selection is done via a lasso drawing tool, allowing the selection of arbitrarily shaped areas of the object surface. The same lasso tool is also used for marking vessel parts, and switching between both modes is accomplished by clicking a tab right above the photo display area. The pattern classes are organized in their *minor groups* (Sec. 3.1.2), and the best-matching pattern groups are listed from top to bottom (Fig. 3 ②). The assignment of a pattern class to the current selection can be made via the

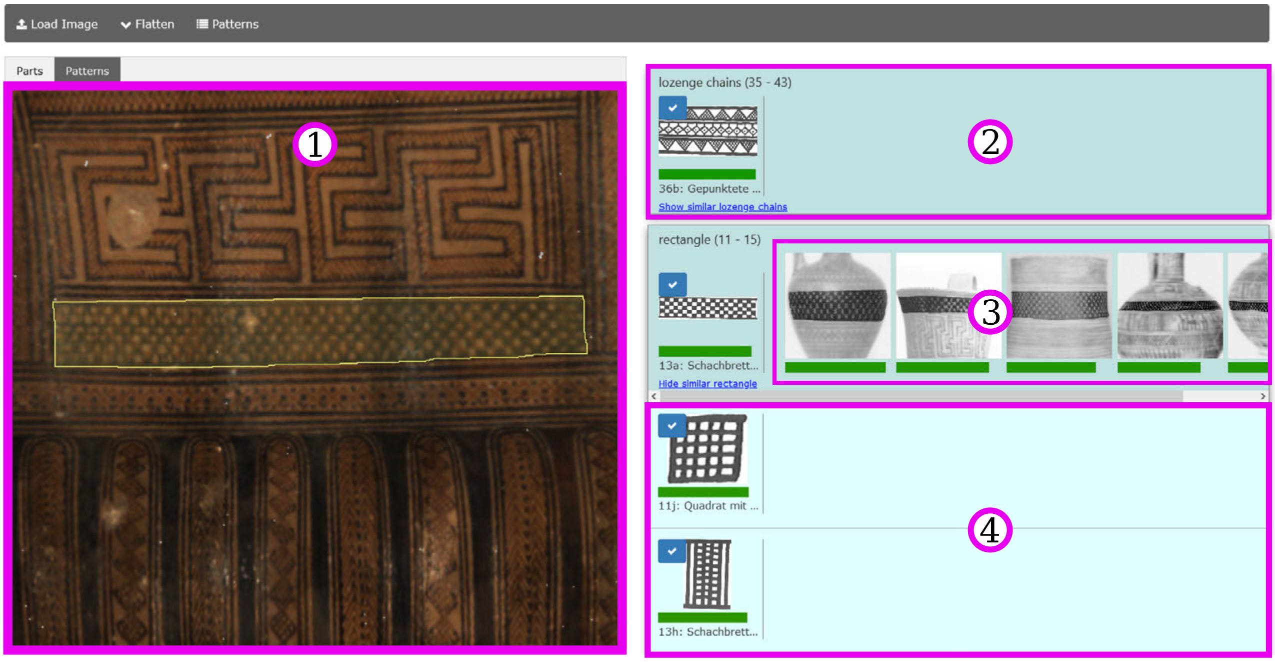


**Figure 3:** Web-based user interface comprised of an annotation area ① with an active selection, and results area (right), with the best matching pattern group ②, the previously annotated inputs for a specific pattern class ③, and further pattern classes belonging to the group of the pattern class directly above it ④. By clicking the respective blue check mark the user can assign the matching pattern class for the current selection.

blue check mark buttons. If annotations for a specific pattern class have been already obtained from other inputs, those annotations are listed next to their sample pattern, sorted (from left to right) according to their similarity to the current selection (Fig. 3 ③). Only the best-matching candidate per group is shown by default, but the other patterns of the group can be displayed on demand (Fig. 3 ④), similar to the nested folder structure display in file browsers.

## 3.4 Extraction of image features

For the similarity computation between a selection and the patterns in our database we need to extract descriptive image features with our CNN. Pottery surface depictions are usually documented as color as well as grayscale photographs. Hence, we apply grayscale conversion [21] to the color photographs to bring the inputs into a common base representation. While this simple preprocessing step suffices for the comparison of photographs with other photographs, we also have to compare photographs to the archetype drawings by Kunisch (Sec. 3.1.2), for pattern classes that do not have any annotations yet. The direct feature extraction from those inputs of a different nature, yields vastly different feature vectors. Consequently, we insert an additional preprocessing step (Sec. 3.4.1) for this case in order to bring the query and reference in a common base representation (sketch representation).

### 3.4.1 Sketch-like input abstraction

The process of transforming the archetype drawings into realistic grayscale images is an ill-posed problem. Hence, we go the other direction and mimic the artistic style used in the documented pattern taxonomy, as is done in a similar fashion by Lee et al. [22]. First, we calculate the discrete Laplacian Δ of the grayscale image. Second, Otsu's method [23] is applied to the inverted Laplace-filtered image, which yields a strictly binary black-and-white image. Otsu's threshold is also applied to the Kunisch pattern images, which are semantically already binary images, but may still exhibit inadvertent intensity variations due to the digitization process. Fig. 4 (from left to right) shows the inputs, intermediate processing steps, and the final outputs calculated by our method for a Kunisch pattern (top) and a sample annotation from a photograph (bottom).

### 3.4.2 Residual neural network

There are many options for selecting a convolutional neural network to extract features from images. We opt for a

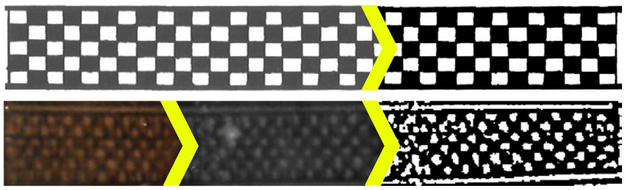

**Figure 4:** Kunisch pattern classes are mapped from grayscale (top left) to black and white (top right). Query selections from input photographs (bottom left) are first converted to grayscale (bottom center) and finally to black and white (bottom right).

well-established architecture that provides high effectiveness and can be easily trained for custom purposes. The residual neural network (ResNet) is an architecture that combines convolutional layers and shortcut connections, making it possible to train very deep models without compromising the convergence of the learning process due to vanishing gradients. We train several networks (ResNet18, ResNet50, ResNet101 and ResNet152, where the number at the end denotes the number of layers) and evaluate the best option for our needs. We use the Pytorch[4] pre-trained networks to perform the fine-tuning. The setup for all the networks is: data augmentation (vertical and horizontal flips), batch size of 32, and Adam optimizer with a learning rate of 0.0001.

We take the Kunisch patterns as input dataset. To make a fair comparison of our implementations, we split the patterns into 70 %/20 %/10 % for training, testing, and validation, respectively. Since all the networks are pre-trained with ImageNet,[5] our goal is to fine tune the networks to learn the patterns in the Kunisch dataset. ImageNet is a large-scale dataset (1.2 M images) with natural images. The pre-training step aims at finding the best latent representations for the images in a classification setup. Even when the network is pre-trained in this large-scale dataset, it is expected that the network has learned to extract meaningful information from the images for the downstream task. Transfer learning is a well-established methodology to take advantage of a pre-trained network to create a specialized neural network for a small, related problem. In our case, the network gains the specialization to represent the patterns via a fine-tuning of the pre-trained ResNet network. Each network is trained (using the same random seed) with the training dataset for classification while the validation set is used for fine-tuning hyper-parameters. Once the networks are trained, we drop the last classification layer from the network,

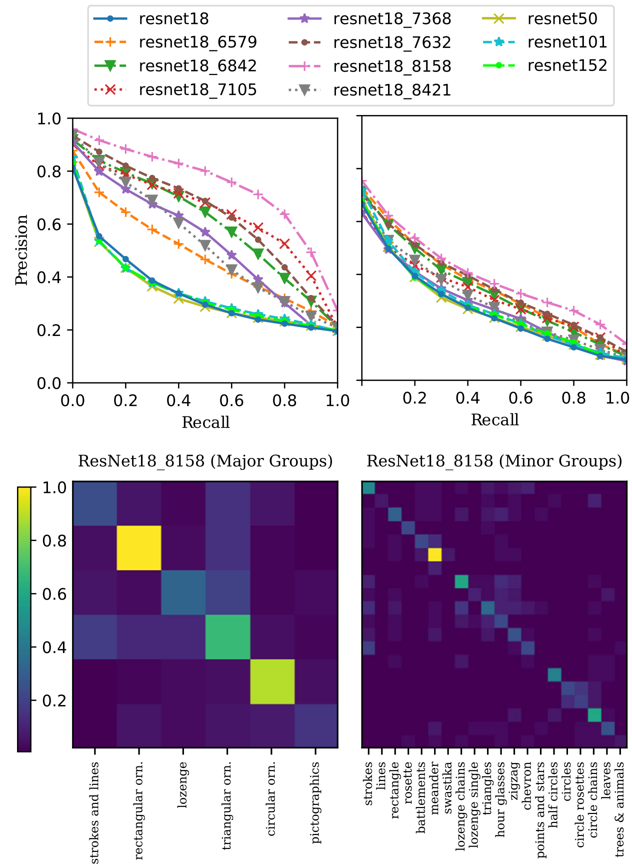


**Figure 5:** The 11-point precision-recall curves for the major- (top left) and minor (top right) pattern groups for different levels of depth of ResNet (18/50/101/152) and different configurations. Bottom left and right show the respective confusion matrices for ResNet18_8158.

and use the output of the global pooling layer as feature vector. Fig. 5, top shows the precision-recall plots for the trained networks, computed over the test set. From this figure we can conclude that 18 layers is the best option to learn from the patterns. It seems deeper networks, such as ResNet50, ResNet101 or ResNet152, require much more data to converge and to generalize, while shallower versions of ResNet would not be able to converge for learning a large-scale dataset and would thus not be useful for transfer learning. After deciding for a certain depth (18 layers) a number of runs of the same ResNet18 base network are conducted as the transfer learning uses several random data augmentations. After trying different configurations, we perform a classification assessment over the validation set. Fig. 5 also shows the 11-point precision-recall curves for other runs of ResNet18, labelled ResNet18_*, where the respective four digit number refers to the network's classification accuracy on the validation set. It can be seen that certain configurations entail a substantial in-

4 https://pytorch.org/
5 https://www.image-net.org/

**Table 1:** Evaluation measures (NN, FT, ST and mAP) for all investigated ResNet architectures. The respective highest values are in bold font.

| Architecture | NN | FT | ST | mAP |
|---|---|---|---|---|
| ResNet18 | 0.6695 | 0.3139 | 0.2536 | 0.3332 |
| ResNet18_6579 | 0.7500 | 0.4376 | 0.3289 | 0.4709 |
| ResNet18_6842 | 0.8506 | 0.5471 | 0.3600 | 0.5971 |
| ResNet18_7105 | 0.8218 | 0.6002 | 0.3784 | 0.6341 |
| ResNet18_7368 | 0.8305 | 0.4838 | 0.3228 | 0.5228 |
| ResNet18_7632 | 0.8822 | 0.5743 | 0.3714 | 0.6297 |
| ResNet18_8158 | **0.9339** | **0.6828** | **0.4127** | **0.7438** |
| ResNet18_8421 | 0.8822 | 0.4688 | 0.3202 | 0.5216 |
| ResNet50 | 0.6753 | 0.2992 | 0.2476 | 0.3247 |
| ResNet101 | 0.7299 | 0.3113 | 0.2556 | 0.3348 |
| ResNet152 | 0.6954 | 0.3097 | 0.2558 | 0.3321 |

crease in performance. It appears that ResNet18_8158 is clearly the best which is also confirmed by established evaluation measures – Nearest Neighbor (NN), First Tier (FT), Second Tier (ST) and mean Average Precision (mAP) – shown in Tab. 1, where this network outperforms all others in every single metric. Fig. 5 bottom shows the confusion matrices for ResNet18_8158, computed from the nearest neighbor pattern in the retrieved list. A high nearest neighbor measure ensures that the most similar pattern is retrieved in the first position during search. It can be seen that the network performs (unsurprisingly) much better on certain pattern groups, such as rectangular or circular patterns, than on rather complex lozenge or pictogram shapes.

Additional annotations resulting from the user-driven documentation process can be leveraged to further fine-tune the network, and since re-training the network is resource expensive, the network is re-trained only once a certain number of new samples are available. This phenomenon is known as concept drift and there are several ways to address it [24].

## 3.5 Similarity computation

To propose accurate pattern class predictions for a surface region selected by the user, we leverage visual similarities, as well the position on the object surface. For the prior we employ the 512-dimensional image features $\{\mathbf{x}\} \subset \mathbb{R}^{512}$ obtained from the ResNet18_8158 network (Sec. 3.4) which we extract for the current selection. To this end, we use the polygon bounding box of the user's lasso selection to crop the annotation surface and mask all image parts outside the selection polygon. The feature vector obtained from this query $\mathbf{x}_q$ can thus be compared to all the pattern class samples provided by Kunisch $\{\mathbf{x}_i\}_{i\in I}$, with $I$ as the index set of pattern classes, as well as all previously annotated patterns $\{\mathbf{x}_j\}_{j\in J_i}$, with $J_i$ as the index set of annotations belonging to the pattern class $i$. We differentiate between feature vectors obtained from the grayscale inputs ($\mathbf{x}^G$) and feature vectors obtained from the binarized (Sec. 3.4.1) inputs ($\mathbf{x}^B$).

We found that the similarity between a query $q$ and a pattern class $i$ is strongly correlated with the similarity of the query to the class's existing annotations, while a weaker correlation was observed between a query and the pattern class samples, as they are artistic representations of a pattern and thus different in nature. Hence, we decide to apply a stronger weighting on the similarities of annotations, if such are present.

We define the similarity of query $q$ to pattern class $i$ as the inverse of the distance

$$d(q,i) = \begin{cases} \frac{\sum_{k=1..|\mathcal{X}_i|} \mathcal{X}_i[k]w(k)}{\sum_{k=1..|\mathcal{X}_i|} w(k)} & \text{if } |J_i| > 0 \\ \theta(\mathbf{x}_q^B, \mathbf{x}_i^B) & \text{otherwise,} \end{cases} \tag{1}$$

with $\theta(\mathbf{x}_1, \mathbf{x}_2) = N^{-\frac{1}{2}}||\mathbf{x}_1 - \mathbf{x}_2||$ as the normalized Euclidean distance, $\mathcal{X}_i = sorted(\{\theta(\mathbf{x}_q^G, \mathbf{x}_j^G)\}_{j\in J_i})$. We apply a weighting to the distances in the sum of Eqn. (1) relative to their appearance in the sorted sequence $\mathcal{X}_i$. The rationale behind it is that the presence of a few very close matches is more important than the overall mean similarity. To this end, we employ a (fast) decaying weighting function $w(k)$ which puts increased emphasis on the best matches. $w(k) = k^{-1}$ is used in all subsequent experiments.

Often the placement of a pattern on the object's surface provides a vital cue of the similarity computation. From the Kunisch taxonomy we have no information regarding the positioning of certain patterns. For the annotations, however, we have such due to the surface part annotation (Sec. 3.1.1). Thus, for each annotation $j$ we can determine a placement histogram $\mathbf{s}_j \in \mathbb{R}^{|\mathcal{V}|}$, where the $k$-th bin $\mathbf{s}_j[k] = |\mathcal{A}_j \cup \mathcal{S}_{\mathcal{V}[k]}|/|\mathcal{A}_j|$, $k \in 1..|\mathcal{V}|$, reflects the annotation's correlation with the $k$-th surface part label (Sec. 3.1.1). $\mathcal{A}_j$ denotes the set of surface pixel of annotation $j$ and $\mathcal{S}_v$ the set of surface pixels of vessel part $v \in \mathcal{V}$. That is, the distance function $\theta$ can be reformulated to $\theta(\mathbf{x}_1, \mathbf{x}_2) = \alpha(1-BC(\mathbf{s}_q, \mathbf{s}_j)) + (1-\alpha)N^{-\frac{1}{2}}||\mathbf{x}_1 - \mathbf{x}_2||$ for the comparision of annotations. $BC$ refers to the Bhattacharyya coefficient [25], which we employ for the distance computation between placement histograms and the weighting factor $\alpha \in [0,1]$ governs the influence of the placement on the similarity computation. The optimal value for $\alpha$ could not be determined within the scope of our experiments as it would require to look at a large amount of annotations.

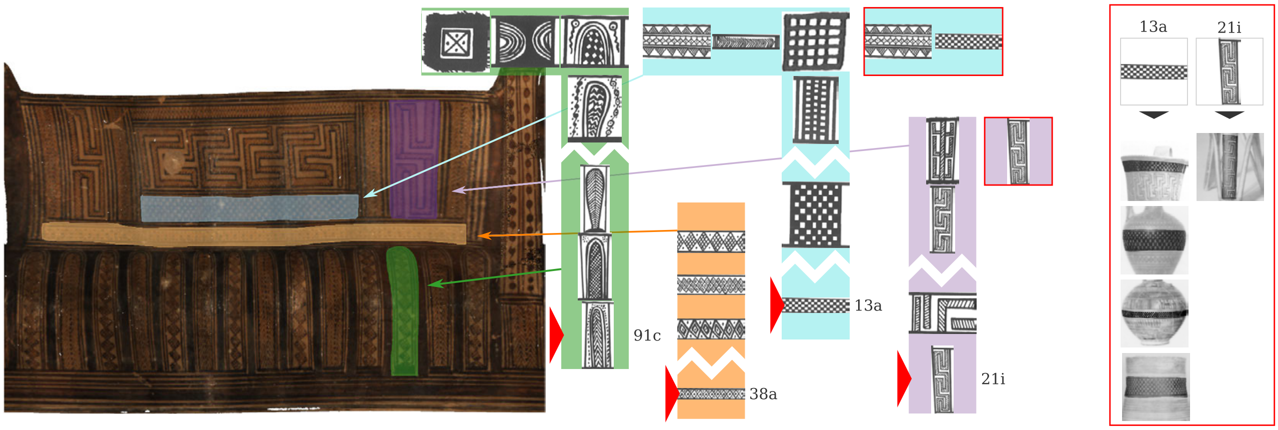


**Figure 6:** Surface annotation of *Vienna KHM IV 1* for four of the object's patterns. The colored panels indicate the group-wise (horizontal) and pattern-wise (vertical) ranking of the correct result according to our automatic recognition. In a second step, the same recognition was tested with a few annotations from other source being present (red box on the right), yielding the rankings marked with a red border.

First experiments, however, indicate that the placement should play a less substantial role, i. e., $\alpha < 0.5$, than the pattern similarity.

## 4 First results

For a first experimental evaluation of our annotation pipeline we investigate four different patterns of varying complexity that are found on the surface of the Attic Geometric pitcher *Vienna KHM IV 1* (Fig. 1). From the unassisted manual annotation effort it is known to which pattern class from the Kunisch taxonomy they can be attributed. For the first scenario we try to identify the correct pattern class based solely on the user selection of the respective surface area and the extracted sample pattern.

The first pattern that we look at is a right-turning hatched meander in vertical panels (according to Kunisch, no. 21i), which appears two times on the left and right side in the neck zone of the pitcher (Fig. 6, left). In the obtained results, the group of meanders is ranked as the closest match, while the exact class within the group is at the fifth position of the best fitting entries. Better yet, the second best entry is actually the same pattern type, except for the turning direction (left instead of right, Fig. 6). The chessboard pattern in a horizontal panel (no. 13a) on the lower neck area of the vessel belongs to the group of rectangular patterns, which is ranked third out of all twenty groups by our retrieval system. Even though the correct pattern class only appears towards the end of the third quarter of the respective group, there are already several good matches among the first few classes (within the first quarter) in the group – c. f. the chessboard pattern at the eighth place (Fig 6). For the lozenge chain with dots within and without in a horizontal panel (no. 38a) at the bottom of the neck zone, the group affiliation was correctly calculated, and the specific lozenge chain at the sides appears already within the first entries of the group. Matching the vertical tongues patterns with stacked M-chevrons and cross-hatched lozenge chains, alternately decorating the belly zone of the vessel, is especially challenging, as the provided taxonomy by Kunisch [1] has no entries that describe the displayed patterns exactly (they show other fillings). According to the domain expert, the depicted patterns correspond to no. 91a to 91g. Nonetheless, these closest equivalents (Fig. 1) are found within the first quarter of all the group instances.

In the second scenario we enriched our initial database with a few real-world artifacts from external sources for several pattern classes. In order to associate the pattern-class depictions on these external vessels, we load the respective vessel photographs into our system, and annotate the surface areas that contain the pattern classes of interest. When we then select the same meander pattern (no. 21i) as in the first scenario, the search result improves substantially. The correct pattern class is now the top result, even after adding just a single annotation of the same meander on a different vessel (Fig. 6, red frames). For the chessboard pattern (no. 13a), four additional occurrences on vessels from different sources are annotated. Again, an improvement is achieved, since the correct group is now listed as the second instead of third best match, and furthermore the associated pattern class is the best match within this group.

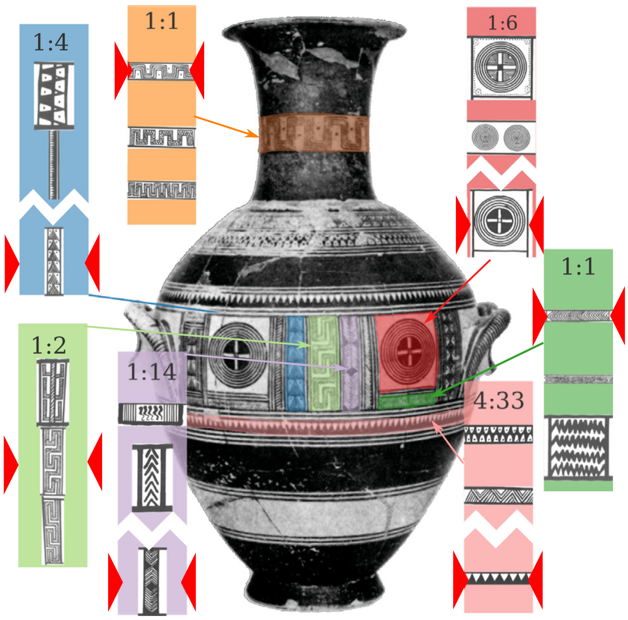


**Figure 7:** Surface annotation of *Athens KER 2146* [26] with the column-separated number at the top of the panels reflecting the group-wise and pattern-wise rank of the correct pattern according to our recognition.

Fig. 7 illustrates the annotation process on another geometric pottery artifact, documented by Kübler [26]. This belly-handled amphora *Athens KER 2146*, decorated in three zones (belly, shoulder and neck), exhibits six clearly distinguishable classes of surface patterns (attribution to Kunisch's taxonomy done again by domain experts): (i) central on the belly a hatched meander shape in a vertical panel (no. 21h); (ii) to the left stacked cross-hatched triangle in a vertical panel (no. 47j); (iii) to the right antithetic chevrons with solid filling in a vertical panel (no. 65f); (iv) further outwards on this belly zone each one large concentric circle enclosing a reserved cross (no. 79d); and (v, vi) below these circles vertical chevrons in horizontal panels, one facing to the right (no. 65d), the other to the left (no. 65e). The whole decorated zone on the belly is framed above and below by ornament bands circumferential the whole vessel with the so-called dogtooth pattern, a row of filled triangles, in our case pointing upwards (no. 53b). Another clearly visible ornament is depicted in the neck zone: a so-called battlement pattern, on this vessel with multiple lines with dots in the open spaces (no. 19b). We annotate those patterns with our system without this domain knowledge and any prior annotations from other sources. Fig. 7 shows the outcome of these annotation efforts, where the colorized vertical panels to the left and right of the vessel visualize the query results for the respective surface area marked with the corresponding color. The colon-separated numbers at the top of each the panels reflect the ranking of the matching search result for the pattern's group and the specific class within the group, respectively, and the two red arrows signify the correct class for the pattern in question. The first four of the six investigated patterns are predicted with a high accuracy, as in these cases the correct group is identified, and furthermore the correct pattern within the group is part of the top results. The fifth pattern – the vertical chevron facing to the left in a horizontal panel (no. 65e) – provides acceptable results: while the correct group is predicted, the correct class within the group is not among the top ten candidates. The sixth pattern – the upwards pointed dogtooth ornament (no. 53b) – is a challenging case, as our recognition system ranks the respective group of zigzag patterns only on forth place, and the correct pattern class within the group is ranked among the classes with a low similarity measure.

## 5 Discussion, limitations and future work

From the scenarios described in Sec. 4 we can conclude that – even without any annotations – the recognition works fairly well for the majority of cases. In some cases, e. g., for the dogtooth ornament (no. 53b), the correct pattern was rated as being very dissimilar, even though the preprocessed query and the reference pattern from Kunisch appear very much alike. The reason why visually similar looking inputs lead to vastly different feature vectors on some occasions could not be established. Yet, it was found that adding annotations generally improved the performance. This is in in line with our expectations, as query and search space entities are of the same modality in that case. At the same time it was observed that a pattern with annotations is not automatically ranked higher than all classes without any annotations, proving that the query preprocessing serves the intended purpose. Note however, that the hand-drawn reference patterns from Kunisch are only used to address the cold-start problem [27], a well known issue with content-based recommender systems. In the long run we plan to replace them completely with annotated patterns taken directly from reference photographs. This way we will have only a single input modality throughout the entire system.

The comparison of the two types of inputs (flattened 3D model and photograph), used for the first results, shows no clear preference for either input type in terms of recognition performance. However, for the practical application the flattened surface has the advantage that it exhibits less perspective distortion and that it represents the continuous object surface as a whole.

Although the annotation was quick and smooth for the shown results, a thorough evaluation of the annotation efficiency with our system remains a future work item. I. e., we intend to assess the usability of our system by conducting an encompassing user study with domain experts.

The second group of limitations is related to the algorithmic capabilities of our applied feature extraction (Sec. 3.4). The employed CNN is an interim solution, due to the small number of available training data, and currently may not be most specific to our domain application. The CNN is not used as a classifier but rather only yields a latent vector for a provided input. This is reflected in a somewhat low robustness with some inputs, since a slight change of the input selection occasionally results in a vastly different ranking. Additionally, inputs of a different nature (photograph and sketch) have to be compared initially, which is addressed by our preprocessing (Sec. 3.4.1), but still poses a challenge for certain inputs. However, contrary to the input related limitations, these algorithmic limitations should be able to be overcome with increasing annotations, which allows the training of a more specialized CNN, or even a pattern class classifier.

**Automatic pre-segmentation of patterns**

At the moment, we require the user to manually encircle patterns to be queried. In the future, we would like to investigate user assistance in this step by proposing a set of pattern candidates that are found in the input images automatically. Even though the topic of image segmentation – the splitting of images into meaningful parts – has been a heavily researched computer vision topic for decades, specialized segmentation tasks (like in our case) still pose a significant challenge. Conventional methods comprise approaches based on thresholding, k-means clustering and watershed algorithms [21], while modern approaches almost exclusively rely on deep learning [28]. However, of-the-self methods are not able to segment the vessel surfaces into their pattern atoms and learning customized segmentation models requires a large amount of training data. Hence, an automatic pre-segmentation of patterns can be feasible once we have collected sufficient data. Yet, even with this potential enhancement, we believe that the user should always have the final choice, such that inadequate automatic segmentation results can always be edited or discarded altogether.

**Pattern recognition**

The pattern recognition part of our pipeline faces two distinct groups of limitations. Firstly, limitations inherent to the input data. Those can be further split into issues related to the artifacts themselves, which result from a poor conservation and include phenomenons like chipped or worn off surface parts and missing parts in general; as well as limitations resulting from the digital data representation. Latter depend on the form of the input data and comprise various acquisition errors, e. g., occlusion, registration errors, etc., in case of 3D digitized objects, and all the challenges characteristic for photographs, resolution, quality, perspective distortions, highlights, etc., in case of images. Another data related limitation specific to our application is that we initially do not have any locality information (which is a vital clue to exclude certain search results) for the pattern classes as this information is not given in the Kunisch taxonomy.

**Kunisch taxonomy**

Another group of limitations stems from basing our pattern classes on the Kunisch taxonomy. This taxonomy was seemingly compiled with a manual annotation process in mind. Hence, it lacks the logical, bottom-up hierarchy, necessary for an automatic classification. E. g., several pairs of pattern classes only differ in their orientation (vertical or horizontal) while rotation-invariance is generally required. A more logical representation of this particular circumstance would be to have just one pattern class for such a pair and (optionally) the orientation as additional flag of an annotation.

In conclusion, the Kunisch taxonomy is a very good starting point for our first prototype, as it constitutes an encompassing structure of geometric patterns with many examples. To the best of our knowledge this is the most encompassing survey of geometric patterns of classic vessels. However, for long term application and larger scale annotation, we want to extend this existing pattern class structure or build up a new taxonomy altogether.

# 6 Applications

An encompassing annotated dataset not only allows for the training of dedicated pattern classifiers by providing labelled image-pattern-pairs, but also enables a variety of other applications.

That is, the structured annotation generated with our system could be leveraged for various search and retrieval applications relevant for archaeological research. Examples include, but are not limited to: searching for objects exhibiting a certain pattern; searching for objects exhibiting a certain combination of patterns; or searching for objects exhibiting a certain pattern at a specific location of their surface.

With dedicated analysis methods deviations in shape or location can be unveiled, leading to new discoveries and insights. Registering the individual occurrences of a certain pattern allows to examine the intra-class variability, identify outliers and even determine a prototypical pattern representation. Based on the location labels, correlations between the pattern classes and surface regions can be revealed.

A structured dataset also allows for visualization and (interactive) exploration. A huge number of methods have been studied to explore multimodal data at all different levels of visual granularity, supporting a broad overview down to close-up. E. g., for the overview visualization of – potentially huge – repositories, visual cluster analysis methods like self organising maps (SOMs) are an established visualization technique. This unsupervised machine learning technique generates a two-dimensional representation of high-dimensional input data. For a simple experiment we apply the approach by Barthel et al. [29] to the pattern drawings from Kunisch and obtained the map depicted in Fig. 8, based on the features obtained from the CNN (Sec. 3.4). Note, that for illustrative purposes, we randomly selected just about half of the entire set of patterns, which were selected at random. The different background colors reflect the different major pattern groups. It can be observed that, even though the CNN is unaware of this coarse classification, very similar clusters are formed where visually similar patterns generally appear close to each other. Other concepts [30, 31] also support the interactive exploration of data and allow the user to investigate selected subsets of the data in more depth and along specific traits in order to reveal correlations.

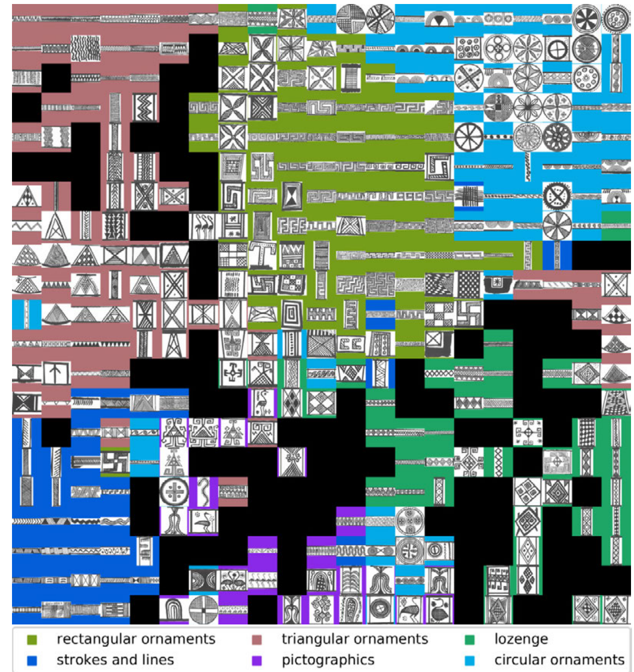


**Figure 8:** Visualization of the Kunisch patterns collection based on feature descriptors obtained from our CNN. The background colors indicate the affiliation major pattern groups.

## 7 Conclusion

We present a semi-automatic annotation pipeline for Geometric patterns on ancient Greek pottery, which greatly simplifies the current purely manual annotation process while requiring no additional effort. The advantage of our proposed process is that the generated annotations exhibit a high degree of structure with references to specific surface areas and digital links to similar patterns. The automatic pattern class prediction based on a selection provides baseline results, and a substantial performance increase is observed even with just a few additional annotations. Although we showed the capabilities of our annotation pipeline with the example of Greek Geometric pottery, due to its availability, we want to stress that the process itself can be generalized for a wide range of other domains.

**Acknowledgment:** We are very grateful to the editors and the reviewers for their valuable comments and suggestions that helped to improve the manuscript.

**Funding:** This article was produced within the interdisciplinary project “Crossmodal Search and Visual Exploration of 3D Cultural Heritage Objects” (P31317-NBL) co-funded by the Austrian Science Fund FWF and the State of Styria, Austria. Benjamin Bustos was funded by ANID – Millennium Science Initiative Program – Code ICN17_002.

## References

1. N. Kunisch. Ornamente geometrischer Vasen. Ein Kompendium. Cologne-Weimar, Böhlau Verlag, 1998.
2. W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu and F.E. Alsaadi. *A survey of deep neural network architectures and their applications*. Neurocomputing 234, 2017, 11–26.

3. E. Beauxis-Aussalet, M. Behrisch, R. Borgo, D. Horng Chau, C. Collins, D. Ebert, M. El-Assady, A. Endert, D. Keim, J. Kohlhammer, D. Oelke, J. Peltonen, M. Riveiro, T. Schreck, H. Strobelt and J. van Wijk. *The Role of Interactive Visualization in Fostering Trust in AI*. IEEE Computer Graphics and Applications 41(6), 2021, 7–12.
4. A. Endert, W. Ribarsky, C. Turkay, W. Wong, I. Nabney, I. Díaz Blanco and F. Rossi. *The State of the Art in Integrating Machine Learning into Visual Analytics*. Comput. Graph. Forum 36(8), 2017, 458–486.
5. M. Chegini, J. Bernard, J. Cui, F. Chegini, A. Sourin, K. Andrews and T. Schreck. *Interactive visual labelling versus active learning: an experimental comparison*. Frontiers of Information Technology & Electronic Engineering 21(4), 2020, 524–535.
6. J. Bernard, M.Zeppelzauer, M. Sedlmair and W. Aigner. *A unified process for visual-interactive labeling*. In: Proceedings of the EuroVis Workshop on Visual Analytics, 2017, 73–77.
7. S. Lengauer, A. Komar, S. Karl, E. Trinkl, I. Sipiran, T. Schreck and R. Preiner. *Semi-automated Annotation of Repetitive Ornaments on 3D Painted Pottery Surfaces*. In: Eurographics Workshop on Graphics and Cultural Heritage. Eurographics Assoc., 2020.
8. J. Zhou, K. Smith, G. Wilsbacher, P. Sagona, D. Reddy and B. Torkian. *Building Science Gateways for Humanities*. In: Practice and Experience in Advanced Research Computing, 2020, 327–332.
9. E.J. Crowley and A. Zisserman. *Of Gods and Goats: Weakly Supervised Learning of Figurative Art*. In: T. Burghardt, D. Damen, W. Mayol-Cuevas and M. Mirmehdi (eds.), Proceedings of the British Machine Vision Conference (Bristol, UK; September 9–13, 2013), 2013, art. 39, 1–11.
10. S. Biasotti, E. Moscoso Thompson, M. Aono, H. Ben, B. Bustos, S. Dong, B. Du, A. Fehri, H. Amin, H. Li, F.A. Limberger, M. Masoumi, M. Rezaei, I. Sipiran, L. Sun, A. Tatsuma, S. Velasco Forero, R.C. Wilson, Y. Wu, J. Zhang, T. Zhao, F. Fornasa and A. Giachetti. *SHREC'17 track: Retrieval of surfaces with similar relief patterns*. In: I. Pratikakis, F. Dupont and M. Ovsjanikov (eds.), 10th Eurographics Workshop on 3D Object Retrieval (Lyon, France; April 23–24, 2017), The Eurographics Association, 2017, 95–103.
11. S. Biasotti, E. Moscoso Thompson, L. Bathe, S. Berretti, A. Giachetti, T. Lejemble, N. Mellado, K. Moustakas, I. Manolas, D. Dimou, C. Tortorici, S. Velasco-Forero, N. Werghi, M. Polig, G. Sorrentino and S. Hermon. *SHREC'18 track: Recognition of geometric patterns over 3D models*. In: A.C. Telea, T. Theoharis and R.C. Veltkamp (eds.), 11th Eurographics Workshop on 3D Object Retrieval (Delft, Netherlands; April 16, 2018), The Eurographics Association, 2018, 71–77.
12. E. Moscoso Thompson, C. Tortorici, N. Werghi, S. Berretti, S. Velasco-Forero and S. Biasotti. *SHREC'18 track: Retrieval of gray patterns depicted on 3D models*. In: A.C. Telea, T. Theoharis and R.C. Veltkamp (eds.), 11th Eurographics Workshop on 3D Object Retrieval (Delft, Netherlands; April 16, 2018), The Eurographics Association, 2018, 63–69.
13. E. Moscoso Thompson, S. Biasotti, A. Giachetti, C. Tortorici, N. Werghi, A. Shaker Obeid, S. Berretti, H.-P. Nguyen-Dinh, M.-Q. Le, H.-D. Nguyen, M.-T. Tran, L. Gigli, S. Velasco-Forero, B. Marcotegui, I. Sipiran, B. Bustos, I. Romanelis, V. Fotis, G. Arvanitis, K. Moustakas, E. Otu, R. Zwiggelaar, D. Hunter, Y. Liu, Y. Arteaga and R. Luxman. *SHREC 2020: Retrieval of digital surfaces with similar geometric reliefs*. Computers & Graphics 91, 2020, 199–218.
14. S. Lengauer, I. Sipiran, R. Preiner, T. Schreck and B. Bustos. *A Benchmark Dataset for Repetitive Pattern Recognition on Textured 3D Surfaces*. In: K. Crane and J. Digne (eds.), Eurographics Symposium on Geometry Processing 2021 (Toronto, Canada; July 12–14, 2021), Computer Graphics Forum 40, 2021, 1–8.
15. C. Steckner. *Begriffliche und empirische Objektordnung*. In: C. Steckner (eds.), Archäologie und Neue Technologien, Schriften des Deutschen Archäologen-Verbandes 10, Deutscher Archäologen-Verband, Freiburg, 1988, 102–139.
16. R. Unnikrishnan and R. Smith. *Combined script and page orientation estimation using the tesseract ocr engine*. In: Proceedings of the international workshop on multilingual OCR, 2009, 1–7.
17. B. Rieck, H. Mara and S. Krömker. *Unwrapping highly-detailed 3d meshes of rotationally symmetric man-made objects*. In: ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences, 2013, 259–264.
18. Y. Liu, B. Zhang and L. Wan. *Automatically unrolling decorations painted on 3D pottery*. In: Proceedings of the Computer Graphics International Conference, 2017, 1–6.
19. R. Preiner, S. Karl, P. Bayer and T. Schreck. *Elastic Flattening of Painted Pottery Surfaces*. In: R. Sablatnig and M. Wimmer (eds.), 16th Eurographics Workshop on Graphics and Cultural Heritage (Vienna, Austria; November 12–15, 2018), The Eurographics Association, 2018, 165–168.
20. P. Houska, S. Lengauer, S. Karl and R. Preiner. *Direct Elastic Unrollings of Painted Pottery Surfaces from Sparse Image Sets*. In: A. Chalmers and V. Hulusic (eds.), 19th Eurographics Workshop on Graphics and Cultural Heritage (virtual, November 4–6, 2021), The Eurographics Association, 2021, 131–140.
21. R.C. Gonzalez and R.E. Woods. Digital Image Processing (3rd Edition), Prentice-Hall, Inc., 2006.
22. Y.J. Lee, C.L. Zitnick and M.F. Cohen. *Shadowdraw: real-time user guidance for freehand drawing*. ACM Transactions on Graphics (TOG) 30(4), 2011, 1–10.
23. N. Otsu. *A Threshold Selection Method from Gray-Level Histograms*. IEEE Transactions on Systems, Man, and Cybernetics 9(1), 1979, 62–66.
24. J. Gama, I. Zliobaite, A. Bifet, M. Pechenizkiy and A. Bouchachia. *A Survey on Concept Drift Adaptation*. ACM Comput. Surv. 46, 2014.
25. A. Bhattacharyya. *On a measure of divergence between two statistical populations defined by their probability distributions*. Bull. Calcutta Math. Soc. 35, 1943, 99–109.
26. K. Kübler. Ergebnisse der Ausgrabungen V 1: Die Nekropole des 10. bis 8. Jahrhunderts. Walter de Gruyter & Co., Berlin, 1954.
27. J. Bobadilla, F. Ortega, A. Hernando and J. Bernal. *A collaborative filtering approach to mitigate the new user cold start problem*. Knowledge-based systems 26, 2012, 225–238.
28. S. Minaee, YY. Boykov, F. Porikli, AJ. Plaza, N. Kehtarnavaz and D. Terzopoulos. *Image segmentation using deep learning: A survey*. IEEE transactions on pattern analysis and machine intelligence, 2021.
29. K.U. Barthel. *Improved image retrieval using automatic image sorting and semi-automatic generation of image semantics*. In: 2008 Ninth International Workshop on Image Analysis for Multimedia Interactive Services. IEEE, 2008.

30. F. Windhager, S. Salisu, R.A. Leite, V.A. Filipov, S. Miksch, G. Schreder and E. Mayr. *Many Views Are Not Enough: Designing for Synoptic Insights in Cultural Collections*. IEEE computer graphics and applications 40(3), 2020, 58–71.
31. S. Lengauer, A. Komar, S. Karl, E. Trinkl, R. Preiner and T. Schreck. *Visual Exploration of Cultural Heritage Collections with Linked Spatiotemporal, Shape and Metadata Views*. In: J. Krüger, M. Niessner and J. Stückler (eds.), 25th International Symposium on Vision, Modeling, and Visualization (Tübingen, Germany, virtual; September 28–October 1, 2020), The Eurographics Association, 2020, 137–144.

## Bionotes

**Dipl.-Ing. Stefan Lengauer**
Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Inffeldgasse 16c, 8010 Graz, Austria
**s.lengauer@cgv.tugraz.at**

Dipl.-Ing. Stefan Lengauer is doing his PhD in the field of cross-modal search and visual exploration of cultural heritage objects. Since 2018 he is a member at Institute of Computer Graphics and Knowledge Visualization and has authored several publications on the topic of shape-based and motif-based retrieval, as well as interactive visualization and reconstruction of ancient pottery.

**Dipl.-Ing. Peter Houska**
Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Inffeldgasse 16c, 8010 Graz, Austria
**p.houska@cgv.tugraz.at**

Dipl.-Ing. Peter Houska is a PhD student at the Institute of Computer Graphics and Knowledge Visualization at Graz University of Technology. His research interests include interactive visualization, geometry processing, and digital archaeology.

**Dr. techn. Reinhold Preiner**
Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Inffeldgasse 16c, 8010 Graz, Austria
**r.preiner@cgv.tugraz.at**

Dr. techn. Reinhold Preiner is a senior researcher in computer science, specialized in 3D computer graphics and visualization. He received his PhD from TU Wien in 2017 and is now with the Institute of Computer Graphics and Knowledge Visualization at Graz University of Technology. His main fields of research include applied geometry processing, rendering, and interactive visualization, involving beside others applications in digital archaeology.

**Dr. phil. Elisabeth Trinkl**
University of Graz, Institute of Classics, Universitätsplatz 3, 8010 Graz, Austria
**elisabeth.trinkl@uni-graz.at**

Dr. phil. Elisabeth Trinkl is a staff member of the Institute of Classics at Graz University. She received her PhD in Classical Archaeology at Vienna University in 1998. Her research interests focus on Greek archaeology in general, the understanding of ancient textiles and the use of computer-assisted methods in archaeological research. She is author of several articles on these topics, a. o. of a volume of the international publication series Corpus Vasorum Antiquorum (CVA), and editor of collections of essays.

**Dr. phil. Stephan Karl**
University of Graz, Institute of Classics, Universitätsplatz 3, 8010 Graz, Austria
**stephan.karl@uni-graz.at**

Dr. phil. Stephan Karl is a scientific member of the Institute of Classics at Graz University. He received his PhD in Classical Archaeology in 2013. His research interests include the study of Early Greek pottery, Roman provincial stone monuments and Roman marble quarries, with a special focus on the application of 3D technologies for visualising and analysing purposes in archaeological research. He is a. o. author of two volumes of the international publication series Corpus Vasorum Antiquorum (CVA).

**Prof. Dr. Ivan Sipiran**
University of Chile, Department of Computer Science, Beauchef 851 3er piso, Santiago, Chile
**isipiran@dcc.uchile.cl**

Prof. Dr. Ivan Sipiran is an Assistant Professor at the Department of Computer Science, University of Chile. He received the PhD in computer science from the University of Chile in 2014. He works in the areas of geometry processing, shape analysis and the application of computer graphics in cultural heritage.

**Prof. Dr. Benjamin Bustos**
University of Chile, Department of Computer Science, Beauchef 851 3er piso, Santiago, Chile
**bebustos@dcc.uchile.cl**

Prof. Dr. Benjamin Bustos is a Full Professor at the Department of Computer Science, University of Chile, and Associate Researcher at the Millennium Institute Foundational Research on Data. He received the Doctoral degree in natural sciences from the University of Konstanz, Germany, in 2006. He leads research projects in the domain of content-based multimedia information retrieval. His research interests include similarity search, 3D object retrieval, multimedia mining, metric/nonmetric indexing, and pattern recognition.

**Prof. Dr. Tobias Schreck**
Graz University of Technology, Institute of Computer Graphics and Knowledge Visualization, Inffeldgasse 16c, 8010 Graz, Austria
**tobias.schreck@cgv.tugraz.at**

Prof. Dr. Tobias Schreck is a Professor and head of the Institute of Computer Graphics and Knowledge Visualization at Graz University of Technology. He previously was Assistant Professor with University of Konstanz, Germany, and Postdoc Fellow with Technical University of Darmstadt. He obtained a PhD in Computer Science in 2006 from the University of Konstanz. Tobias Schreck works in the areas of Visual Analytics, Information Visualization, and Applied 3D Object Retrieval.