

MillefioriAnalyzer: machine learning, computer vision and visual analytics for provenance research of Ancient Roman artefacts

Conference or Workshop Item

Published Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Open Access

Wiebel, A., Gloger, O. and Eckardt, H. ORCID:

<https://orcid.org/0000-0001-9288-5624> (2025)

MillefioriAnalyzer: machine learning, computer vision and visual analytics for provenance research of Ancient Roman artefacts. In: 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, 26-28 Feb 2025, Porto, Portugal, pp. 807-814. doi: <https://doi.org/10.5220/0013098900003912> (ISBN: 9789897587283) Available at <https://centaur.reading.ac.uk/120108/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.5220/0013098900003912>

Publisher: IVAPP

copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

MillefioriAnalyzer: Machine Learning, Computer Vision and Visual Analytics for Provenance Research of Ancient Roman Artefacts

Alexander Wiebel¹^a, Oliver Gloger¹^b and Hella Eckardt²^c

¹*UX-Vis Research Group, Worms University of Applied Sciences, Worms, Germany*

²*Department of Archaeology, University of Reading, Reading, U.K.*

{*wiebel, gloger*}@*hs-worms.de*, *h.eckardt@reading.ac.uk*

Keywords: Visual Analytics, Computer Vision, Digital Humanities, Ancient Roman Artefacts, Archaeology, Millefiori.

Abstract: In this position paper, we explore ways to digitally support provenance research of ancient Roman artefacts decorated with millefiori. In particular, we discuss experiments applying visual analytics, computer vision and machine learning approaches to analyze the relations between images of individual millefiori slices called florets. We start by applying automatic image analysis approaches to the florets and discover that image quality and the small overall number of images pose serious challenges to these approaches. To address these challenges, we bring human intuition and pattern recognition abilities back into the analysis loop by developing and employing visual analytics techniques. We achieve a convenient analysis workflow for the archaeologists by integrating all approaches into a single interactive software tool which we call *MillefioriAnalyzer*. The software is tailored to fit the needs of the archaeological application case and links the automatic image analysis approaches with the interactive visual analytics views. As appropriate for a research software, *MillefioriAnalyzer* is open-source and publicly available. First results include an automatic approximate ordering of florets and a visual analytics module improving upon the current manual image layout for further analytic reasoning.

1 INTRODUCTION

A distinctive group of Roman copper alloy objects has millefiori ('a thousand flowers') decoration, small polychrome patterns created by arranging slices (florets) of glass rods (canes) in elaborate, highly symmetrical and very striking patterns on a copper-alloy base. An estimated 1200 of these objects are found across the Roman world but existing overview studies are long out-dated (e.g. (Exner, 1939); (Henry, 1933)). Despite nearly a hundred years of study, we have not progressed beyond a general assumption of production in the Rhineland, Belgium or the Danube provinces on general stylistic grounds. Research is inhibited by the fact that the material is dispersed, currently very poorly documented and the complexity of the decorative motifs. Each stud or brooch with millefiori decoration can have between three and five zones of decoration, using multiple different motifs, with as many as 100 florets decorating a single object.

This pilot project explores whether it is possible

to use machine learning (ML), computer vision (CV) and, in particular, visual analytics (VA) to understand millefiori designs. In particular, and as a first step, we ask whether it is possible to identify florets from the same cane – and therefore potentially identify objects made in the same workshop. These florets are very small (3x3mm), and distorted from the process of stretching the cane, making comparison with human eyes without any digital support difficult. In a second step, we intend to identify the floret designs, explore the 'grammar' of the motifs and color combinations and study the placement of different design on the jewellery. This could help to understand if specific combinations occur in particular regions. Thus, the main contributions of this application paper are:

- An analysis of the challenges of applying CV and ML techniques to millefiori images.
- A prototypical software, *MillefioriAnalyzer*, which represents a first step towards supporting the analysis of collections of millefiori images by combining CV, ML and visualization techniques into an application-specific VA system.
- Automation of parts of the archaeological research workflow for millefiori research.

^a <https://orcid.org/0000-0002-6583-3092>

^b <https://orcid.org/0000-0002-0791-4273>

^c <https://orcid.org/0000-0001-9288-5624>



Figure 1: **Top:** Example of a copper alloy and enamel stud inlaid with millefiori, from Chepstow, UK (left) and Usk, UK (right). Left: © The Trustees of the British Museum. Shared under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) licence. Right: © Amgueddfa Cymru - Museum. **Bottom:** Illustration of a cane (left) and florets (right).

- Perspectives for follow-up research steps.

The paper is structured to describe and discuss the contributions in the order they are mentioned above.

2 BACKGROUND AND DATA

We chose two very similar studs (figure 1), from Chepstow now in the British Museum and from Usk, now in Caerleon Museum (belonging to *Amgueddfa Cymru – Museum*) for our experiments. They are virtually identical in size (51mm) and design (spirals of pure white against a blue background (Bateson D1), squares of the three by three white and blue chequerboard pattern within a red frame (Bateson A7), flowers with eight white petals and a white centre surrounded by a red circle against a blue background (Bateson C13) with a central panel filled by alternating three by three white and blue chequerboard pattern within a red frame (A7), and five by five white and blue chequerboard pattern within a blue frame. The two studs were found ca. 25 km apart in the 19th century ((Brailsford, 1954), 56, pl. XXI, No. 6; (Henry, 1933) fig. 41.1; cf. https://www.britishmuseum.org/collection/object/H_1891-0327-12 or https://www.britishmuseum.org/collection/object/H_1891-0327-9; (Lee, 1862), 56, pl. XXVIII, No. 14). It seems likely that they were made

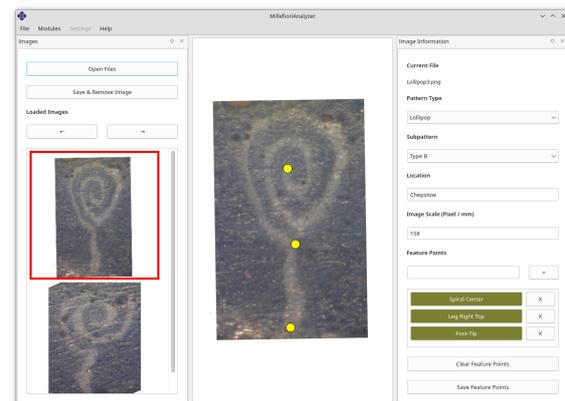


Figure 2: Image annotation module. Markers for three feature points have been added/edited in the image and meta-data have been edited/provided in the side panel.

in the same workshop, but the question addressed here is whether it is possible to strengthen that suggestion by identifying florets from the same cane.

In our experiments we process images (photos) focusing on one floret each (see figures 2 and 4). Many of these images are prone to high-frequency reflections (see e.g. figure 2) and artefacts resulting from damage to the material of the florets (see figure 5).

3 RELATED WORKS

Previous applications of machine learning in archaeology have focused on highly standardized objects such as coins (Deligio et al., 2023; Kiourt and Evangelidis, 2021; Aslan et al., 2020; Schlag and Arandjelovic, 2017) while applications to more variable forms of material culture such as pottery were perhaps less successful (van Helden et al., 2022). Bickler (Bickler, 2021) highlights how small archaeological datasets with complicated contextual information and poor surface images can be problematic, all issues that affected this project.

The visual analytics part of this project was motivated by exactly these issues. We are not the first to approach archaeology problems in this way. Employing visual analytics in archaeological research has already been suggested over a decade ago (Llobera, 2011). Knowledge discovery in rock art research (Deufemia et al., 2014) is one example where different petroglyphs have been presented in a way similar to the cluster view in our system. Another system exhibiting a clustering view as core part for the analysis of pottery motifs has just recently been presented by Li et al. (Li et al., 2024). The research questions and thus the clustering and supporting interaction and visualization techniques in these applica-

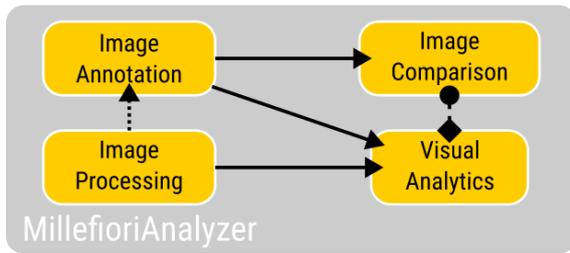


Figure 3: MillefioriAnalyzer system architecture with edges indicating data flow and calls between different modules.

tion areas are different to those in millefiori research and thus in the MillefioriAnalyzer.

4 REQUIREMENTS

We conducted repeated preliminary interviews with an archaeologist and her students to understand current practices, identify deficiencies, and determine our design requirements. The archaeologist is from a university and is our collaborator. Based on the interviews we created two personas (Lidwell et al., 2010) for typical users: an archaeology student and an archaeology professor. For these personas we derived user stories which then were used to design and implement the different features of the system.

Currently, matching and comparison of florets is done by eyeballing. Due to the large amount of florets belonging to one cane, this process is tedious and did not lead to results with the desired quality. Thus, automatic matching of similar florets or technical support for comparison and grouping of florets is required.

In order to explore the *grammar* of motifs and color combinations, the archaeologists currently manually arrange the florets on slides in a presentation software. In this process all additional (meta) information is lost in the sense that it is not directly accessible from the slide. Thus iteratively refining the arrangement using new information, e.g. from algorithmic image analysis, becomes cumbersome if not impossible. An interactive layout and annotation tool retaining the connection to the original data and its meta information is needed.

5 SYSTEM ARCHITECTURE

MillefioriAnalyzer is a research software. It serves as a tool in archaeology research and as prototype for developing new computational approaches for the archaeology research. Thus the system comprises multiple modules (see figure 3) exploring different ways

to analyze the millefiori images:

1. *Image processing module* (see section 6), incorporating CV and ML techniques
2. *Image annotation module* (see figure 2), for manually adding further meta data
3. *Visual analytics module* (see figure 7 and section 7.1), for interactive overall analysis of floret relationships
4. *Image comparison module* (see figure 8 and section 7.2), for detailed analysis of potentially which florets were originally neighbors in a cane.

This allows the archaeologists use and evaluate multiple analysis approaches in one tool with a uniform interaction philosophy and look-and-feel across all modules. Additionally, the integration of the different approaches allows to link them for an overview first and details on demand approach (Shneiderman, 2003). A connection (see figure 3) between the visual analytics module, for overview, and the image comparison module, for inspecting differences between individual images, is an example for such a link.

The system has been implemented in Python using PyQt as general GUI framework, it is platform-independent, and it is freely (GNU LGPL v3) available at <https://gitlab.rlp.net/ux-vis-public/millefiorianalyzer> and <https://doi.org/10.5281/zenodo.14589448>.

6 COMPUTER VISION AND MACHINE LEARNING

The intention of the computer vision and machine learning part of our interdisciplinary work is the determination of millefiori images that were originally produced from the same cane. This task requires us not only to figure out the similarity between different millefiori images but also to determine the correct order of the millefiori images in their original canes. Furthermore, the machine learning algorithms must differentiate between different canes to assign millefiori images correctly to the cane they belong to.

A major drawback of the data in our project is the lack of training images as we could only gather 174 original millefiori images for our purposes. This small amount of training examples can be a serious disadvantage to learn significant patterns for image similarity with machine learning methods and with deep learning architectures in particular. Yet another difficulty is the fact that, as mentioned above, images of the millefiori slices frequently show damage (see figure 5) and inclusions; there are also strong reflections that makes the identification of the originally

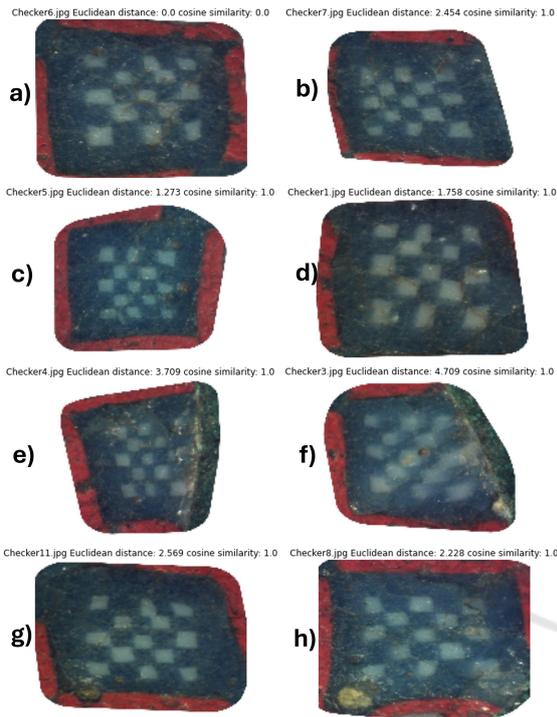


Figure 4: A reconstructed cane with an image sequence a)-h) for a cosine similarity threshold of 0.9.

composed patterns even more difficult. Hence, the fully automatic determination of image similarity including the correct assignments to their original canes is a great challenge for our project.

6.1 Automatic Image Processing

The challenging steps require the automatic calculation of the similarity of two images. There exist several methods to compute image similarities like SSIM (structural similar index measure) or pixel-based RSME (root mean square error), normalized cross correlation, keypoint-detectors like SIFT (scale-invariant feature transform) or SURF (speeded up robust features) features and others. Due to the high variability in image appearance in combination with the low number of images for ML approaches, we started by using classical image processing techniques. These were implemented in the image processing module (figure 3).

Applying these techniques, we observed that significant reflections lead to inappropriate SIFT-keypoints and classical techniques like SSIM values were not able to capture the image similarity between the millefiori patterns. Normalized cross correlation produced the best results of all classical image processing methods that we tested. It often showed higher values for similar images. However, this kind

of template matching results are not sufficient to determine image similarities as they do not take the displacement of the most significant millefiori patterns inside the cane into consideration.

6.2 Machine Learning Approaches

There are frequently used deep learning methods that can learn latent feature spaces of the images. For instance, (variational) autoencoders (Ballard, 1987; Kingma, 2013) or Siamese networks (Koch et al., 2015) learn to generate feature spaces, into which images can be transformed. Distances between images embedded in feature space can then be calculated. They represent image similarity measures. Siamese networks with contrastive or triplet loss provide more potential to learn image similarities than (variational) autoencoders, since the millefiori images can be combined to reach a larger amount of training examples. In case of contrastive loss the training examples are pairwise combined for training example generation. Siamese networks that are steered by minimization of triplet loss require a triplet training example consisting of a positive, negative and an anchor image yielding even more training examples. Hence, in contrast to autoencoder and variational autoencoders Siamese neural architectures offer the advantage to collect much more training examples in order to achieve better training results for image similarity.

We applied autoencoders, variational autoencoders and Siamese networks using both contrastive loss and triplet loss. Since there exist no supervised similarity values between the millefiori images, no quantitative results can be provided to evaluate the calculated image similarities. However, we observed the most promising results for Siamese networks using triplet loss as they assign similar images to one cane. For Siamese network training using triplet loss, we divided the millefiori images into 5 categories whereupon each category represents the same underlying millefiori pattern. We iterate through all millefiori images and determine a new anchor image in each iteration. Positive examples are taken from all other images of the same category and negative examples are sampled from the other four categories. Thus, each triplet example consist of an anchor image, a positive and a negative image. We applied transfer learning and used a ResNet50 (He et al., 2016) as backbone with pre-trained weights, which were trained for images of the ImageNet database.

As mentioned, the aim is find a sequence of florets which come from the same cane. For testing purposes, we compute the embeddings of all images in feature space. We start with the embedding vector of



Figure 5: Challenging image quality due to damaged florets: missing parts, small holes, scratches or scraped top layer.

a randomly chosen millefiori image (actual predecessor) and determine that image as successor in the cane that has the highest cosine similarity between the two embedding vectors. The successor then assumes the role of the new predecessor, for that the new successor is determined in the same manner. If no successor with a higher similarity as a pre-defined threshold is found, we start with a new cane. However, we must take into account, that the original cane may not be complete due to missing intermediate millefiori images. The results (see figure 4) show that canes can be reconstructed but depend highly on the used cosine similarity threshold. If we choose too low thresholds, then the reconstructed canes might include millefiori images from another category. If too high thresholds are used, then many millefiori images do not get a successor and the reconstructed canes might be too small. This shows that improvements to the techniques are needed, but that an automatic separation of different categories of floret designs is possible.

6.3 Manual Image Processing

The results obtained by the approaches discussed in the previous subsection indicate that it might not be feasible to process the original images directly and automatically. This insight is supported by another experiment where we applied a simple edge detection filter to the gray scale version (figure 6 bottom left) of one of the original images (figure 6 top left). The result is shown in the top row of figure 6. Obviously the relevant features, the checker pattern, are hidden in a plethora of edges resulting from image noise (e.g. reflections). Histogram equalization before edge detection improves the visibility of the pattern a bit but also strengthens the surrounding noise. Both of these edge images are not useful.

An edge image that could be helpful for comparing florets should look like the bottom right image in figure 6. This image, however, has been obtained by first manually tweaking a value for brightness thresholding, three times erosion and dilation afterwards in order to remove smaller artefacts, and finally applying the edge detection to this manually improved image. Such a manual process, however, would be too time consuming when applied to many images and our attempts to automate it have not been successful yet.

7 VISUAL ANALYTICS

Motivated by the challenges encountered in the experiments using computer vision and machine learning described in section 6, we decided to bring humans back into the analysis loop. Using visual analytics approaches (Wong and Thomas, 2004) we can combine human intuition and pattern recognition abilities with the science of mathematical deduction, in our case computer vision and ML, to derive connections between the images and insight from these connections.

7.1 Visual Analytics Module

The goal of this module is to integrate the original image data, the results from the automatic analysis and a human expert into the most effective analysis process possible. To achieve this multiple different perspectives on the images and the metadata are needed. Currently these perspectives are provided by three linked visual analytics views as shown in figure 7. Based on the analysis of the requirements (section 4) and on the continued use of the system by the experts, additional views will probably be added in the future. The three

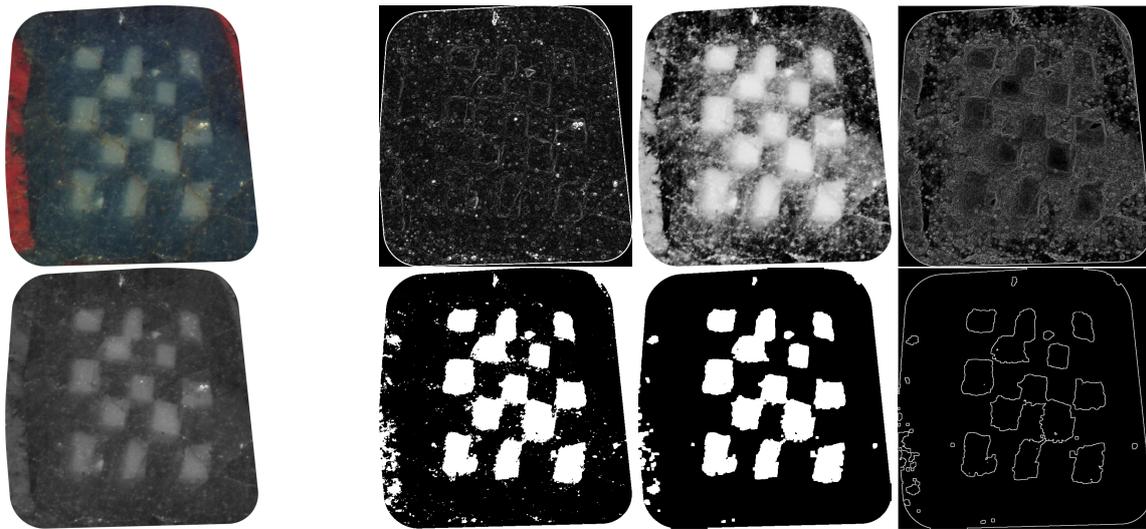


Figure 6: A checkerboard floret processed in different ways shows challenges resulting from the image quality. **Left column:** color image, grayscale image (basis for all other steps). **Right top row:** Sobel edge detection, equalized, Sobel on equalized. **Right bottom row:** thresholding with value 80, previous image eroded and dilated three times, Sobel on previous image.

existing views are described in the following. As all views are linked, highlighting a floret in one of the views will also highlight it in all other views.

7.1.1 Table View

The most basic, but also the most verbose of the current views, is the table view (top left in figure 7). It lists all loaded images row-wise. Each row contains the filename of the floret image, the basic pattern type present in the floret, a subpattern type, the location where the object decorated by the floret has been found, the number of feature points in the meta data, and a thumbnail of the floret. The table can be sorted by each of the columns. Each floret can be selected and deselected (for highlighting). The corresponding row will be highlighted accordingly.

7.1.2 Parallel Coordinates View

The parallel coordinates view (bottom left in figure 7) presents axes for all properties (metadata) of the florets. The name and the thumbnail are not shown as they do not belong to the metadata. As is characteristic for parallel coordinates plots (Inselberg and Dimsdale, 1990), each data point, in this case each floret, is represented by a line connecting the locations on the different axes which correspond to values of its properties. Lines and thus florets can be highlighted by clicking on the lines or the line crossings at the axes. The highlighting state of lines at the clicked location will be toggled accordingly. Highlighted florets are rendered as solid lines, others as dashed lines.

7.1.3 Image Clustering View

The image clustering view (right in figure 7) is intended to support semantic layout and grouping of florets for analysis purposes. Currently, the view supports to drag images from the repository row at the bottom to the main canvas above it. In the main canvas, images can be dragged to move them to the desired position. Related and thus closely positioned images can be grouped by frames and groups can be annotated by text. As this view, again, is linked to all other views, images highlighted in other views are highlighted by frames in both, the repository row and the main canvas. Like all other views, the cluster view can be used to highlight florets for interactive analysis. This is possible in the canvas and the repository.

In the current version the clustering view only supports manual layout and grouping. For the future it is planned to provide initial layout using dimension reduction and projection techniques like t-SNE (Van der Maaten and Hinton, 2008) as well as initial grouping from the implicit clustering of such an algorithm. Furthermore the clustering currently only serves direct visualization and illustration purposes. In the future the groupings will be made available as floret properties to the other views. Thus, the groups can be the basis for further interactive analysis.

7.1.4 Module Summary

The visual analytics module in its current form serves as a replacement for the currently practiced manual layout on presentation slides. The layout produced in the clustering view, the annotation and the currently

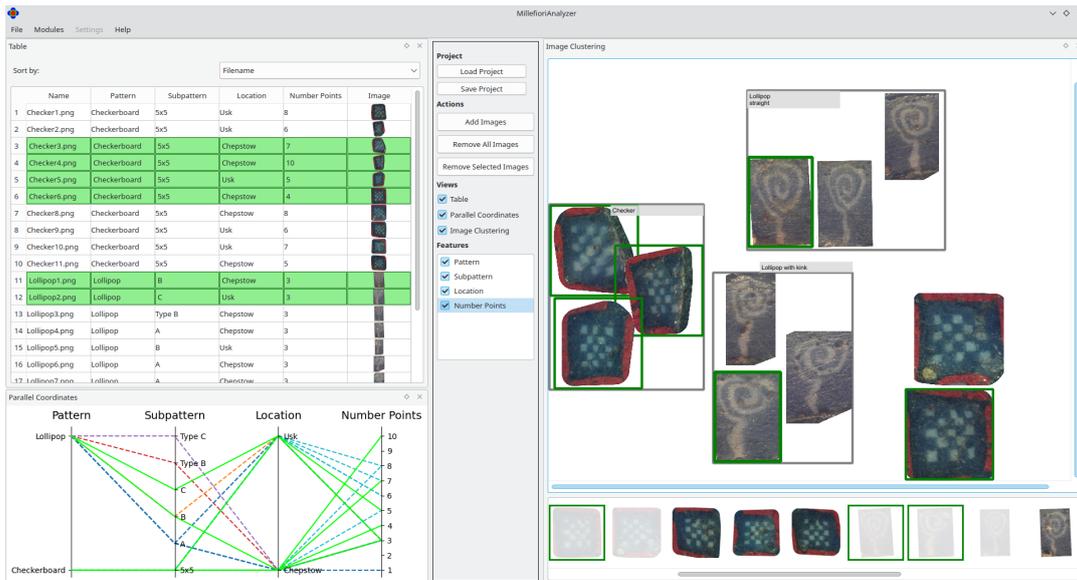


Figure 7: Visual analytics module. Three views are shown: table, parallel coordinates and image clustering. Several florets are highlighted. Three semantic groups have been formed manually in the clustering view..



Figure 8: Image comparison module. Two images have been superimposed for visual comparison.

highlighted florets can be saved in a *project file* and thus the connections between the layout and the floret meta data can be retained. Such connections are not preserved in a suitable way for automatically reading and pre-processing when using slide presentations. Restoring the analysis status can serve presentation purposes and allows to continue with the interactive analysis whenever needed. While the employed visual analytics techniques might appear to quite basic, they were chosen specifically to suit the application.

7.2 Image Comparison Module

The image comparison module enables users to visually check how similar two florets are and in which feature they differ. For this purpose the module over-

lays the two images as shown in figure 8. The user can decide which of the images will be shown on top. The comparison is possible by toggling the visibility of the top image. This works well because humans are good in recognizing differences when images are shown in immediate succession (Healey and Enns, 2012). To study the differences in more detail, the module allows users to adapt the opacity of the images in order to see both images at the same time (see figure 8).

The images are initially aligned according to key points stemming from the annotation module. The user interface allows to rotate images around the main keypoint to manually improve the alignment. The size of the images is determined based on an image scale attribute storing *pixels/mm*. The size is important because the patterns in consecutive florets of a cane are expected to have very similar size. This module is integrated into the analysis workflow by allowing users to choose two images in the *image clustering view* and directly open the comparison.

8 CONCLUSION AND FUTURE WORK

In this position paper we introduced the archaeological application case of millefiori provenance research and described the challenging nature of the image data arising in this context. We developed the research software *MillefioriAnalyzer* which addresses the requirements of the archaeologists by integrating and combining the automatic and interactive approaches

explored in this paper. The current automatic approaches, which are intended to find floret images stemming from the same millefiori cane, allow for a distinction between the different floret types. The precision of the ordering in the computed image sequence needs further improvement. The visual analytics part of *MillefioriAnalyzer* allows for interactive layout of the florets for visual analysis and retaining the connection to the meta data at the same time. The connection to the meta data has been lost in the archaeologists previous layout workflow.

As described throughout this paper, the *MillefioriAnalyzer* software and the archaeological analysis are not yet complete. As usual in new digital humanities projects (Jänicke, 2016), more iterations of development and evaluation are needed. We will incorporate manually segmented images into our machine learning part. These segmentations, provided by archaeological experts, contain the most significant patterns of the florets and thus will positively influence the minimization of the Siamese cost function. As a result, the calculated image similarities will be less confounded by damage, reflections and other noisy patterns. In the future we will also acquire photographs of higher quality and of more millefiori artefacts.

ACKNOWLEDGMENTS

We would like to thank T. Wendler, P. Kretler, M. Kretler and S. Brender for their implementation support, and S. Lambert-Gates and S. Sarkar for supporting us in preparing the photographs and tracings.

REFERENCES

- Aslan, S., Vascon, S., and Pelillo, M. (2020). Two sides of the same coin: Improved ancient coin classification using graph transduction games. *Pattern Recognition Letters*, 131:158–165.
- Ballard, D. H. (1987). Modular learning in neural networks. In *Proceedings of the sixth National conference on Artificial intelligence-Volume 1*, pages 279–284.
- Bickler, S. H. (2021). Machine learning arrives in archaeology. *Advances in Archaeolog. Practice*, 9(2):186–191.
- Brailsford, J. W. (1954). *Guide to the antiquities of Roman Britain*. The Trustees of the British Museum.
- Deligio, C., Tolle, K., and Wigg-Wolf, D. (2023). Supporting the analysis of a large coin hoard with AI-based methods. In *CAA2023 Conf. Proc.* CAA, Zenodo.
- Deufemia, V., Indelli Pisano, V., Paolino, L., and de Roberto, P. (2014). A visual analytics system for supporting rock art knowledge discovery. In *Computational Science and Its Applications-ICCSA 2014: 14th Internat. Conf.*, pages 466–480. Springer.
- Exner, K. (1939). Die provinzialrömischen Emailfibeln der Rheinlande. *Bericht der römisch-germanischen Kommission*, pages 31–121.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE CVPR*, pages 770–778.
- Healey, C. and Enns, J. (2012). Attention and visual memory in visualization and computer graphics. *IEEE TVCG*, 18(7):1170–1188.
- Henry, P. (1933). Émailleurs d’occident. *Prehistoire II*, pages 66–146.
- Inselberg, A. and Dimsdale, B. (1990). Parallel coordinates: a tool for visualizing multi-dimensional geometry. In *Proceedings of the first IEEE conference on visualization: visualization90*, pages 361–378. IEEE.
- Jänicke, S. (2016). Valuable research for visualization and digital humanities: A balancing act. In *Worksh. on Vis. for the Digital Humanities, IEEE VIS*, volume 7.
- Kingma, D. P. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kiourt, C. and Evangelidis, V. (2021). Ancoins: Image-based automated identification of ancient coins through transfer learning approaches. In *International Conf. on Pattern Recognition*, pages 54–67. Springer.
- Koch, G., Zemel, R., Salakhutdinov, R., et al. (2015). Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2, pages 1–30. Lille.
- Lee, J. E. (1862). *Isca Silurum. Or an illustrated catalogue of the Museum of Antiquities at Caerleon*. Longman, Green, Longmans & Roberts.
- Li, J., Lai, C., Zhang, H., and Yuan, X. (2024). PM-Vis: A visual analytics system for tracing and analyzing the evolution of pottery motifs. *IEEE TVCG*, 30(6):3022–3034.
- Lidwell, W., Holden, K., and Butler, J. (2010). *Universal principles of design, revised and updated: 125 ways to enhance usability, influence perception, increase appeal, make better design decisions, and teach through design*. Rockport Pub.
- Llobera, M. (2011). Archaeological visualization: towards an archaeological information science (aisc). *Journal of Archaeological Method and Theory*, 18:193–223.
- Schlag, I. and Arandjelovic, O. (2017). Ancient roman coin recognition in the wild using deep learning based recognition of artistically depicted face profiles. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 2898–2906.
- Shneiderman, B. (2003). The eyes have it: A task by data type taxonomy for information visualizations. In *The craft of inform. vis.*, pages 364–371. Elsevier.
- Van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-sne. *J. of machine learning research*, 9(11).
- van Helden, D., Mirkes, E., Tyukin, I., and Allison, P. (2022). The arch-i-scan project: Artificial intelligence and 3d simulation for developing new approaches to roman foodways. *J. of Comp. Applic. in Archaeology*.
- Wong, P. C. and Thomas, J. (2004). Visual analytics. *IEEE Computer Graphics and Applications*, 24(5):20–21.