Can Stones be Made (Artificially) Intelligent? A Pilot Study for AI-based Repeated Pattern Detection of Architectural Decoration in Roman Asia Minor

Julie Verlinden[1,2], Thorsten Mahieu[2], and Toon Goedemé[2]

[1]Ancient History Research Group, Department of History, Ghent University, Belgium
[2]PSI-EAVISE Research Group, Department of Electrical Engineering, KU Leuven, Belgium

The disciplines of History, Archaeology and Computer Vision have each stimulated the development and use of diverse digital tools within their own domain. But what if we dare to ask a research question positioned right at the crossroad of these fields, necessitating the computer-based analysis of people's past via both inscriptions and monuments? In this paper we return to the dawn of our era (1st century BCE - 3rd CE) in the historical region of Asia Minor (roughly modern-day Turkey) to put this to the test. Our starting point? The many legacy photographs of so called *Bauornamented* buildings. With the help of pattern recognition techniques, we attempted to design a (proof-of-concept) AI-algorithm capable of detecting and segmenting the various complex designs in pictures of such embellished building blocks. In doing so, we focused on mitigating the bottlenecks of the traditional time-consuming *manual* comparative approaches, allowing to unlock this data's true historical potential[1].

Keywords: Computer Vision, Repeated Pattern Detection, AI, Architectural Decoration, Roman Asia Minor

1 Introduction

Our starting point is a simple, yet significant observation about the Roman empire: archaeologists and ancient historians encounter a remarkable degree of similarity in the material culture and architecture of its many studied provincial cities (Van Oyen and Pitts, 2017). Focusing on buildings, we observe that, despite differences of detail,

[1] This research was funded by the Research Foundation Flanders (FWO), grant number 11PDX24N.

constructions look remarkably similar, also in Asia Minor. What processes explain this similarity in an empire where the central government was bureaucratically minimalist (Garnsey and Saller, 2015, 1987) and local communities enjoyed a lot of autonomy (Zuiderhoek, 2009)? Furthermore, the socio-cultural effects of imperial politics and policies were never top-down, nor straightforward, and could take many forms (Mattingly, 2011). A pivotal dataset, omnipresent but so far overlooked in any attempt to unravel this commonality[2], forms our dataset of choice: that of architectural decoration. Also known as *Bauornamentik*, the canon of geometrical and vegetal designs found on a building has no structural function, yet is present on nearly every monument with some grandeur in the empire. And, it is prone to fashions, changing over time and space, yet displaying the above-mentioned similarities (Ginouvès and Martin, 1985; Vandeput, 1997). Especially in Asia Minor, many monuments are also linked to surviving inscriptions. These provide information on who commissions or pays for these buildings; who sets trends or copies them; who plays which role in the game of social status. These texts, in other words, connect the buildings and their decorations to the people behind the stones: the elite protagonists who built their lives in these ornamented cities. In brief, we investigate if the similarities identified in the ornaments match with other (un)known forms of connectivity (institutional, cultural, social, religious) between the same cities and their elites, known from written sources[3].

The opportunity for computer vision techniques here pertains to the quantity and granularity of detail that needs to be captured. At this stage, our data compendium contains just under 2,000 images of decorated blocks from ca. 300 monuments on 40 archaeological sites. Since every building block typically exhibits multiple different designs, mapping and comparing these manually is practically impossible, both in terms of time consumption and because the sheer number of research components exceeds the human capacity of analysis and comparison. Hence, AI-based automatic pattern recognition (step 1) and comparison (step 2) techniques are indispensable, placing architectural decoration studies on an exciting and innovative new footing. Indeed our proposed workflow consists of a two-step approach: 1) isolate the separate repetitive decorative patterns as identified in literature from the legacy photographs,[4], yielding a digital library of pattern snippets per different canonized design class, 2) sub-cluster the latter per ornament class to unveil, measure and quantify the mentioned (dis)similarity via an *embedding learning* or *metric learning* approach. This consists of

---

[2] Different scholars have attempted to answer this question, e.g. based on the concepts of globalization (Pitts and Versluys, 2015) or on social and cultural revolutions (Wallace-Hadrill, 2008). To date, however, the main merit of the field of architectural decoration studies (typically part of the broader domain of *Bauforschung*) lies in its contribution to establishing chronological frameworks for archaeological buildings, sites and regions e.g. Gliwitzky (2010). As a result, this archaeological dataset's potential to answer broader socio-historical questions has so far remained underexplored.

[3] Roman Imperial Asia Minor provides an ideal and relevant laboratory to perform such stylistic, architectural, urban and historical exercise, considering the comparatively high degree of urbanization and monumentalization of the vast peninsula (Willet, 2020), the number of extensively excavated ancient cities in the region (Gates, 2011), their well-studied architectural (Yegül and Favro, 2019)(esp. ch.10) and exceptionally rich epigraphical record. Moreover, there is the simultaneous academic recognition of both ancient historians and archaeologists of the significance of these topics (Koller et al., 2021; Lohner-Urban and Quatember, 2020; Marek, 2016)

[4] For a comprehensive overview of the nomenclature of the main decorative designs, see eg.Vandeput (1997) in English, Cavalier (2005) in French, Rumscheid (1994) in German. This Ph.D. project also works on the disambiguation of the existing nomenclature as apt terminology in English is for some patterns still problematic. This will be presented in the database which will be made publicly available.

converting a photo/snippet into a descriptor vector (*the embedding*) that describes the image content. Distances between those vectors are then trained to correspond to the visual (dis)similarity which can vary per region and/or time window. This allows to compare a set of snippets of the same family/class of patterns (e.g. Ionic kyma decoration) and divide them into subsets of (more) homogeneous groups (e.g. egg-and-dart, egg-and-tongue, intermediate forms) according to the needed granularity (De Feyter et al., 2023).

In this paper, we focus on the first step of our workflow, the application of artificial intelligence (AI) for the automated recognition of the architectural decoration canon in Roman Asia Minor. Our core task is thus automating the identification and classification of repetitive decorative patterns found in Roman architecture, thereby enabling a more efficient exploration and deeper understanding of cultural transmission and stylistic evolution across the Roman Empire.

2 Proposed AI pipeline

In short, we have developed an AI-driven system that can recognize, segment, and classify these architectural patterns from photographic data. In this paper, we explore state-of-the-art computer vision techniques that would render this task totally automatic, without any of the manual labor traditionally involved in analyzing these motifs. Indeed, the very fact that these patterns consist of a unit that consistently repeats multiple times on the image makes a clue to detect and isolate each pattern without supervision.

A limited but representative dataset of 50 images of already published Roman blocks from across Asia Minor was compiled, serving as the basis for training, testing and refining our algorithm. Since this dataset purely served the purpose of establishing our novel pipeline, we did not impose any further chronological or spatial selection of the Roman building blocks depicted. Preprocessing techniques were applied to improve the accuracy. First, Contrast Limited Adaptive Histogram Equalization (CLAHE) was used (Reza, 2004), where histogram equalization is applied to small parts of the image ('tiles') and combined via bilinear interpolation. This improved the contrast between shadow and sunlight without disrupting major contrasts. In addition, gamma correction further refined the contrast. For the actual repreating pattern detection, we explored and compared two types of algorithms, leading to the new fusion of them presented in this paper: semantic pattern discovery to segment each pattern from the image and convolutional neural networks (CNNs) activation analysis for detecting the recurring pattern units. In the text below, we will first tackle each separately, to then demonstrate what our novel synthesis provides as assets and improvements.

2.1 Repeating pattern segmentation

After having scanned the existing repeating pattern detection algorithms to identify the most apt one for the task here at hand (Das and Mandal, 2018; Mahmudnejad et al., 2023; Qu et al., 2020; Sakai et al., 2024), a first algorithm that showed potential was "Unsupervised Semantic Discovery through visual patterns detection" by Pelosin et al. (2021). It is the most optimized state-of-the-art open source algorithm that can detect multiple patterns by deploying the Canny algorithm, DAISY and SLIC techniques. The algorithm works independently of a fixed geometric scheme and proved more performant than GRASP (Liu and Liu, 2013). This makes it possible to detect a wider

Figure 1: Repeating pattern segmentation results by using Pelosin et al. (2021) on a typical image from our dataset, here from the ancient city of Attaleia (figure after Kaymak (2009) figs.142-147 & 270, adaptations by the authors).

range of diverse patterns, even if they vary in scale, orientation or (partial) visibility.

As can be seen on Figure 1, after having finetuned the hyperparameters through a limited labeled ground truth set (13 images), the algorithm was able to successfully segment few patterns, but some inconsistencies remained, e.g. two separate patterns were segmented as one (the pink coloured segment on Figure 1, 4th coloured band from below combines a bead-and-reel string with Ionic kyma decoration into a single segment), some motifs remained undetected (see Figure 1, the wide floral scroll band located in the middle of the image, probably due to the lack of strict identical repetitions in the displayed part), while other plain sections of the building block were identified as repeated motif (see Figure 1, the orange coloured segment). Interestingly, the inscription on the building block was also detected (in violet). To be clear, the goal of our approach here is not to be able to individually detect every of the composing units/letters for an inscription (as will be presented in the second step in subsection 2.2 for all other patterns). This would be redundant work as all inscriptions on legacy blocks have already been identified, translated and made accessible via corpora and collections such as *Inschriften griechischer Städte aus Kleinasien* (IGSK), *Monumenta Asiae Minoris Antiqua* (MAMA), *Tituli Asiae Minoris* (TAM), Regional Epigraphic Catalogues of Asia Minor (RECAM) and the PHI Greek Inscriptions Database (https://inscriptions.packhum.org/). What we do want to achieve is (as happened here) that when present on a block, the inscription is segmented so that archaeologists/historians are notified of its presence (even when badly visible/weathered) and can incorporate the translation and content into their analysis.

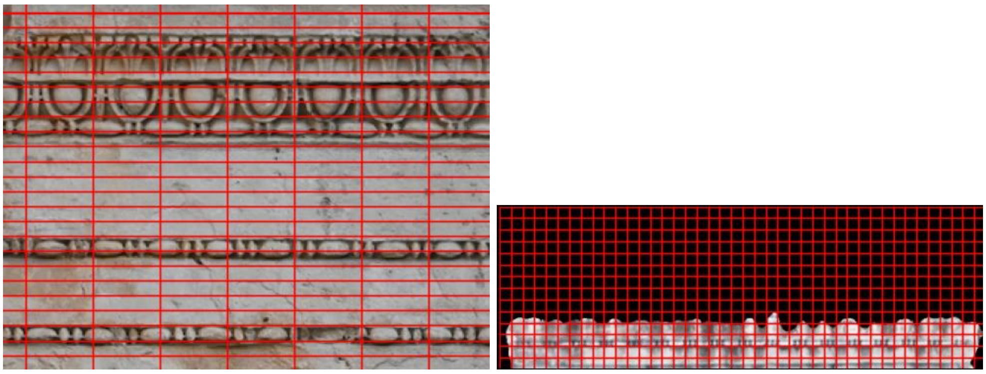

Figure 2: Results of Lettry et al. (2017). Left: on an image of a building block with multiple repeating patterns (figure by the authors). Right: on a wavely cut-out part (figure after Başaran and Ergürer (2018) fig.15a, adaptations by the authors).

2.2 Repeating unit detection

The second algorithm "Repeated Pattern Detection Using CNN Activations" (Lettry et al., 2017) was initially not considered because it is unable to detect multiple patterns at the same time. Due to the need to outline the various units for classification, however, this was also tested. The hypothesis we made was that the most prominent pattern would be detected. After that, the other patterns would be covered. As shows Figure 2, however, this algorithm was not consistent in outlining the repetitive unit. It pinned the horizontal distribution to the second pattern starting from below. However, the vertical distribution did not consistently grip on the same part of the unit of that pattern. Therefore, we sought to give each pattern separately as input to frame the individual units.

2.3 Combined approach

Since both algorithms individually did not meet our goals, we propose to combine them into a novel pipeline, providing a more robust system capable of handling most complexities of ancient decorative motifs. We deployed the technique of Pelosin et al. (2021) to segment and extract the patterns present in the image first. Next, for each of the cut-out pattern segments, we used the method of Lettry et al. (2017) to frame the individual repetitive units. This combined approach is illustrated in Figure 3. The first algorithm was performed with one of the best hyperparameter combinations, derived from the CSV file, where each separate strip of detected patterns was stored as separate segments. We then used these segments as input in the second algorithm to frame the individual figures.

It was important that the image used did not contain other repetitive elements that would derive the algorithm from the main pattern. This problem arose when the segmentation was not straightforward and wavy or jagged lines were present at the edges, in which the algorithm found regularity (see Figure 2, right). Therefore, we adjusted segmentation to take the maximum x and y coordinates of the segment and cut out a rectangle. Lastly, the algorithm sometimes drew unneeded horizontal lines because it recognized multiple patterns on a vertical plane while there was only one. Therefore, we adjusted the code so that it only drew vertical lines and a frame around the full horizontal pattern (see Figure 3, the right green arrow).

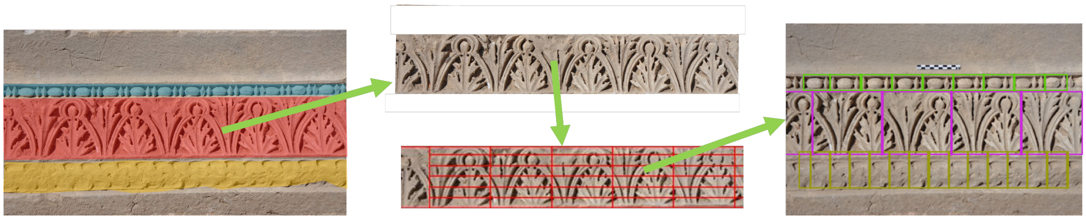

Figure 3: Our combined pipeline demonstrated on a block from the ancient city of Sagalassos. Each of the detected pattern segments is cut-out from the image and individually analysed to find the repeating unit, before composing again together (figure after Poblome et al. (2022) p.39, adaptations by the authors).

3 Results

Put all together, our newly fused algorithm consists of the following pipeline: after image cleaning, we extract all pattern segments from an image by applying a modified version of the initial semantic discovery algorithm with optimized parameters. Subsequently, the bounding boxes for each segment are determined via the CNN filter activations. Using the stored coordinates, the segment is then provided with a frame in a random color in the right place in the image. Within this frame, the vertical separations between the figures are drawn based on the bounding boxes. If no coordinates are stored, they are determined via template matching. The bounding boxes are also stored for future use. This process is repeated for each segment, after which the entire figure becomes visible. We tested this process on both the labeled training data and new test images from the dataset.[5] As our research is rooted in and feeds a visual-comparative inquiry, not a statistical one, we opt to present an evaluation of our detection based on the comparison of the frequency of the unit separations with the repetition frequency of the pattern (see Figures 4 to 7).

First, our algorithm was successfully able to find and segment all horizontal patterns, except for one. In Figure 4, 11 horizontal pattern segments are visible but the so-called anthemion frieze (3rd from above) remained undetected. Interestingly, the algorithm did find the same anthemion pattern in Figure 7 where it also occurred (top moulding, yellow bounding boxes). Looking at the nature of the motif in Figure 4, the lack of more than one full repeating unit, the damage and the fact that this pattern class typically displays in-class variation (both mentioned segments in Figures 4 and 7 are classified in literature as anthemion friezes but the way in which the individual flowers are sculpted and combined into a repeating sequence can vary greatly) might explain this difference in detection. In one case, the algorithm also identified two discrete patterns as one horizontal segment, *in casu* a bead-and-reel string surmounted by Ionic kyma decoration (egg-and-tongue) (see pink bounding boxes in Figure 6). In Figure 4, however, the same positioning of these two designs was also present (even twice), but here both were correctly split up into two separate horizontal segments (see light green and brown bounding boxes for the first and second design counted from the top as well as the red and olive coloured units for the third and fourth motif starting from the bottom of the image).

Second, turning to the detection of the individual repeating units, three observations are to be noted. In few occasions, we encountered the issue of oversegmentation (see

[5] Our code is openly available on `https://github.com/duster3000/pattern-detection-roman-architecture`.

green and blue bounding boxes in Figure 6) and undersegmentation of the repetitive unit (see the bead-and-reel string at the bottom of Figure 4 with the purple bounding boxes). In addition, we observed that the algorithm did not always consistently follow one and the same fixed start and endpoint for all attestations of patterns of one class. For example, the delineation of an individual bounding box for the egg-and-dart motif at times cuts right through the onset of an egg design (see Figure 5, green bounding boxes) while in Figure 6 for example, the individual repetitive unit detected by the algorithm for the same motif starts at a different point in the sequence (the pink bounding boxes fully enclose the egg). Lastly, the delineation of the bounding boxes once intersected with important semantic information in a motif, for example when cutting right through the core and central vein of an acanthus leaf (see Figure 5, yellow bounding boxes). These last two remarks touch upon the core of the second step of our workflow and analysis described above, the sub-clustering per ornament class to unveil, measure and quantify the mentioned (dis)similarity via an *embedding learning* or *metric learning approach*. Comparing two non-identically delineated repeating units of the same pattern would cause a false bias. The same holds for comparing two half acanthus leaves instead of the full leaf.

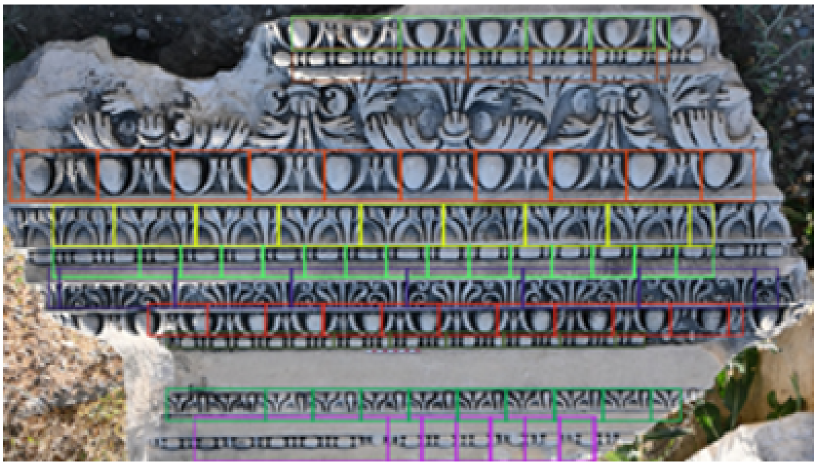

Figure 4: Resulting pattern detections on a decorated block from the ancient city of Myra (figure by the authors).

4 Discussion

With this novel synthesis method, we demonstrated that the system effectively detects and segments common horizontal patterns, such as the bead-and-reel or egg-and-dart motifs. Remarkably, this process runs totally autonomously, without any human intervention. However, we also encountered some limitations in terms of the consistency and definition of the unique repeating unit. Further research could therefore refine the model to accommodate more complex patterns (eg. sloping or oblique ones) and handle incomplete or degraded architectural elements.

At the same time, we realized that the unique start and end point of such a repeating unit is actually not defined. As there are multiple possibilities per pattern, setting the unit of analysis is not unambiguous but the golden rule is simply: consistency. Based on the lessons we have learned from developing this pipeline as well as looking into our future historical analysis' component, we actually doubt it is necessary to each

Figure 5: Resulting pattern detections on a decorated block from the ancient city of Aizanoi (figure after Jes et al. (2010) fig.66, adaptations by the authors).

time find the repeating unit. Rather, it is the full sequence and its interconnections that hold most semantic information. Therefore, in the future we will also invest in and orient us more towards *semantic segmentation* rather than *object detection* (eg. YOLO) to automatically retrieve all decorative patterns.

5 Conclusion

Based on its rich attestation in Roman Imperial Asia Minor, we are assembling and studying the region's *Bauornamentik* corpus from a multidisciplinary and methodologically innovative perspective. In doing so, we aim to for the first time connect this hitherto overlooked archaeological dataset to the broader socio-historical study of Roman urban culture to generate important new insights into the development of provincial Roman urbanism. However, as manually comparing all legacy photographs is unfeasible, we presented in this paper a proof-of-concept AI-algorithm that could serve as a first methodological step to detect the different repeating motifs. To do so, we fused two existing algorithms into a novel pipeline, capable of detecting the repetitive unit of each horizontal pattern segment.

As we have presented here, the basis for classification has been reached, but there is still room for improvement in the following areas: ensuring a consistent start and end point for all motifs, addressing both over-segmentation and under-segmentation as well as combining patterns or splitting patterns appropriately. In the coming time, we thus aim to extend our approach to handle non-horizontal patterns, expand the labeled dataset, and foremost, to improve the detection of certain patterns that are currently not always being recognized (e.g. the scroll pattern in Figure 1). Based on the insights from this research, we will likewise shift our focus more towards *semantic segmentation* instead of object detection to automatically capture all decorative patterns.

As demonstrated in this paper, the integration of AI with traditional historical methods opens up new possibilities for large-scale cultural analysis and cross-site comparisons, providing novel insights into the socio-historical diffusion of people, artistic styles and the networks of Roman artisans. Our next challenge encompasses the

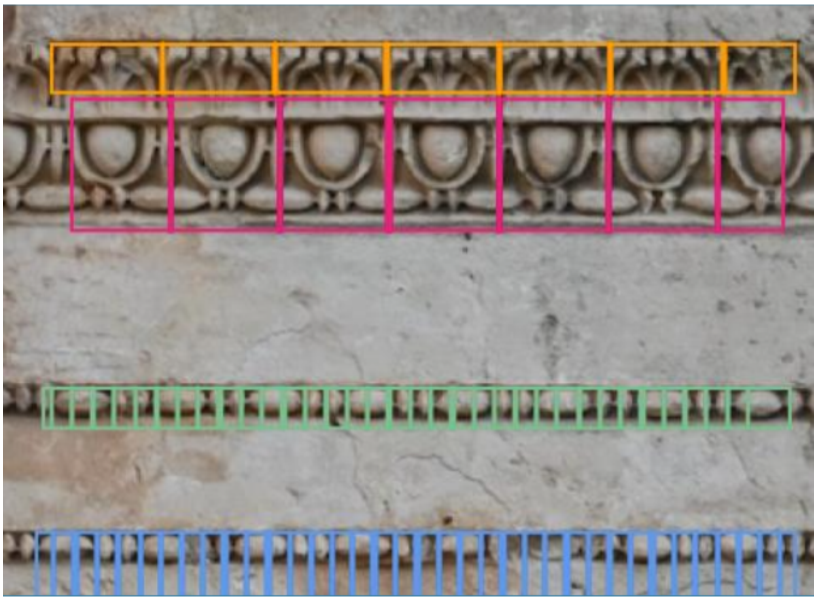

Figure 6: Resulting pattern detections on a decorated block from the ancient city of Attaleia (figure by the authors).

Figure 7: Resulting pattern detections on a decorated block from the ancient city of Letoon (figure by the authors).

development of step 2, a technique to measure fine-grained similarities and differences between the patterns detected, in order to be able to pinpoint and compare subtle stylistic details across archaeological sites. Only by using high-performance embeddings (e.g. SimCLR(Chen et al., 2020)) we will unveil subtle differences between patterns within a class, offering the crucial yet previously invisible clues for our novel historical analysis. In this way, we argue, ornamental patterns of (dis)similarity within and between ancient cities can be discerned on an unprecedented scale, as well as the structure of the communities and social networks responsible for these buildings, whose actions might explain the sameness in the decoration.

References

Cevat Başaran and H. Ertuğ Ergürer, editors. *Roman theater of Parion: excavations, architecture and finds from 2006 - 2015 campaigns*. İÇDAŞ, 2018.

Laurence Cavalier. *Architecture romaine d'Asie Mineure: Les monuments de Xanthos et leur ornementation*. Ausonius Éditions, 2005. doi: 10.4000/books.ausonius.3237.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

Sugata Das and Sekhar Mandal. Figure spotting in indian heritage image. *Journal of Cultural Heritage*, 32:133–143, 2018.

Floris De Feyter, Kristof Van Beeck, and Toon Goedemé. Deep visual recognition for the real world. Phd dissertation, 2023.

P. Garnsey and R. Saller. *The Roman Empire: Economy, Society and Culture*. University of California Press, 2015.

P. Garnsey and R.P. Saller. *The Roman Empire: Economy, Society and Culture*. University of California Press, 1987.

Charles Gates. *Ancient Cities: The Archaeology of Urban Life in the Ancient Near East and Egypt, Greece, and Rome*. Classical studies and archaeology. Routledge, 2011.

René Ginouvès and Roland Martin. *Dictionnaire méthodique de l'architecture grecque et romaine. 1 Matériaux, techniques de construction, techniques et formes du décor.* Collection de l'École française de Rome 84,1. Ecole française de Rome, Rome, 1985.

Christian Gliwitzky. *Späte Blüte in Side und Perge: Die pamphylische Bauornamentik des 3. Jahrhunderts n. Chr.* Peter Lang Verlag, Lausanne, Switzerland, 2010. doi: 10.3726/978-3-0351-0128-7.

Kai Jes, Richard Posamentir, and Michael Wörrle. Der tempel des zeus in aizanoi und seine datierung. In Klaus Rheidt, editor, *Aizanoi und Anatolien: neue Entdeckungen zur Geschichte und Archäologie im Hochland des westlichen Kleinasien*, pages 59–87. 2010.

Gamze Kaymak. *Die Cumanın Camii in Antalya: ihre Baugeschichte und ihre byzantinischen Ursprünge: Bauaufnahme, Bauforschung, Denkmalpflege = Antalya Cumanın Camii: mimari tarih ve Bizans kökeni: rölöve, yapı analizi, anıt koruma ve bakımı*. Number 9 in Adalya Supplementary Series. Suna-Inan Kıraç Akdeniz Medeniyetleri Araştirma Enstitüsü, 2009.

Karin Koller, Ursula Quatember, and Elisabeth Trinkl, editors. *Stein auf Stein*. Number 9 in Keryx. Unipress Verlag, 2021.

Louis Lettry, Michal Perdoch, Kenneth Vanhoey, and Luc Van Gool. Repeated pattern detection using cnn activations. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 47–55. IEEE, 2017.

Jingchen Liu and Yanxi Liu. Grasp recurring patterns from a single view. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2003–2010, 2013.

Ute Lohner-Urban and Ursula Quatember. *Zwischen Bruch und Kontinuitat: Architektur in Kleinasien am Ubergang vom Hellenismus zur römischen Kaiserzeit*. Byzas 25. Ege Yayınları, Istanbul, 2020.

A. Mahmudnejad, E. Andaroodi, and M. Saadatseresht. Advanced clustering of architectural geometric ornaments using small scale machine learning, case study of ilkhanid geometric patterns. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-4/W1-2022:417–422, 2023. doi: 10.5194/isprs-annals-X-4-W1-2022-417-2023.

Christian Marek. *In the Land of a Thousand Gods: A History of Asia Minor in the Ancient World*. Princeton University Press, 2016.

D.J. Mattingly. *Imperialism, Power, and Identity: Experiencing the Roman Empire*, volume 1 of *Miriam S. Balmuth Lectures in Ancient History and Archaeology*. Princeton University Press, 2011.

Francesco Pelosin, Andrea Gasparetto, Andrea Albarelli, and Andrea Torsello. Unsupervised semantic discovery through visual patterns detection. In *Structural, Syntactic, and Statistical Pattern Recognition: Joint IAPR International Workshops, S+SSPR 2020, Padua, Italy, January 21–22, 2021, Proceedings*, pages 272–281. Springer, 2021.

M. Pitts and M.J. Versluys. *Globalisation and the Roman World: World History, Connectivity and Material Culture*. Cambridge University Press, 2015.

Jeroen Poblome, Peter Talloen, Inge Uytterhoeven, and Julie Verlinden. *Spolia, stenen met een verhaal*. Peeters; Leuven, 2022.

Hong Qu, Yanghong Zhou, K. P. Chau, and P. Y. Mok. Repeated pattern extraction with knowledge-based attention and semantic embeddings. Proceedings of the 14th IADIS International Conference Computer Graphics, Visualization, Computer Vision and Image Processing 2020, CGVCVIP 2020 and Proceedings of the 5th IADIS International Conference Big Data Analytics, Data Mining and Computational Intelligence 2020, BigDaCI 2020 and Proceedings of the 9th IADIS International Conference Theory and Practice in Modern Computing 2020, TPMC 2020 - Part of the 14th Multi Conference on Computer Science and Information Systems, MCCSIS 2020, pages 99–106. IADIS Press, 2020.

Ali M Reza. Realization of the contrast limited adaptive histogram equalization (clahe) for real-time image enhancement. *Journal of VLSI signal processing systems for signal, image and video technology*, 38:35–44, 2004.

Frank. Rumscheid. *Untersuchungen zur kleinasiatischen Bauornamentik des Hellenismus*. Number Bd. 14 in Beiträge zur Erschliessung hellenistischer und kaiserzeitlicher Skulptur und Architektur. P. von Zabern, 1994.

Masato Sakai, Akihisa Sakurai, Siyuan Lu, Jorge Olano, Conrad M. Albrecht, Hendrik F. Hamann, and Marcus Freitag. Ai-accelerated nazca survey nearly doubles the number of known figurative geoglyphs and sheds light on their purpose. *Proceedings of the National Academy of Sciences*, 121(40):e2407652121, 2024.

A. Van Oyen and M. Pitts. *Materialising Roman Histories*. University of Cambridge Museum of Classical Archaeology Monographs. Oxbow Books, 2017.

Lutgarde Vandeput. *The architectural decoration in Roman Asia Minor: Sagalassos : a case study.* Studies in Eastern Mediterranean archaeology 1. Brepols, Turnhout, 1997.

A. Wallace-Hadrill. *Rome's Cultural Revolution*. Cambridge University Press, 2008.

Rinse Willet. *The geography of urbanism in Roman Asia Minor*. Equinox Publishing Ltd., Sheffield, South Yorkshire, 2020.

F.K. Yegül and D.G. Favro. *Roman Architecture and Urbanism: From the Origins to Late Antiquity*. Cambridge University Press, 2019.

Arjan Zuiderhoek. *The Politics of Munificence in the Roman Empire: Citizens, Elites and Benefactors in Asia Minor*. Cambridge University Press, 2009.