



Original article

Novel approaches for enhanced visualisation and recognition of rock carvings at Stonehenge

Gavin Leong^{a,*}, Matthew Brolly^a, Hugo Anderson-Whymark^b, David J. Nash^{c,d}, Jon Bedford^e^a Centre for Environment and Society, University of Brighton, Brighton, United Kingdom^b Scottish History and Archaeology, National Museums Scotland, Chambers Street, Edinburgh, United Kingdom^c School of Applied Sciences, University of Brighton, Brighton, United Kingdom^d School of Geography, Archaeology and Environmental Studies, University of the Witwatersrand, Johannesburg, South Africa^e Historic England, York, United Kingdom

ARTICLE INFO

Article history:

Received 11 December 2024

Accepted 8 July 2025

Available online 25 July 2025

Keywords:

Difference of Gaussians

Pseudo-depth map

Meshnet

Rock carving

Carving enhancement

Machine learning

ABSTRACT

The sarsen uprights at Stonehenge feature the largest panels of Early Bronze Age axe-head carvings in the world. Archaeologists use these carvings to better understand the significance of the monument. Between 2011 and 2012, the analysis of laser scanning and photogrammetric data led to the identification of 71 axe-head carvings and one dagger carving, in addition to the 44 carvings already known. Recent advances in carving visualisation and machine learning warrants a reanalysis of this data using new methods. Two novel techniques for carving visualisation, difference of Gaussians and pseudo-depth mapping, are introduced and compared to four recent techniques, radiance scaling, openness, distance between meshes, and extended difference of Gaussians. On the northwest face of Stone 53, difference of Gaussians highlighted the presence of two previously unidentified carvings, ten potential areas of carving, and nine alternative interpretations on previously found carvings. Pseudo-depth mapping revealed the presence of a further two previously unidentified carvings. In addition, an existing classifier for 3-D shape representation, MeshNet, is converted into a technique for carving recognition. MeshNet achieved 90.7 % accuracy on labelling samples of surfaces at Stonehenge with and without carvings, close to the benchmark performance of 91.9 % on ModelNet40. Both difference of Gaussians and pseudo-depth mapping can be implemented for visualisation of highly faded rock carvings in under two hours and under ten minutes respectively, while the application of MeshNet serves as a feasibility study of semi-automated carving recognition.

© 2025 The Author(s). Published by Elsevier Masson SAS. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Until the late 20th century, Britain was thought to lack the rich prehistoric rock art heritage found in continental Europe, especially figurative rock carvings [1]. However, advances in remote sensing technology have revealed that this perception is partly due to the highly eroded nature of these carvings, which make them challenging to identify [2]. This is especially evident at Stonehenge, where analysis by Abbott and Anderson-Whymark of laser-scans produced in 2011 revealed 72 Early Bronze Age axe-head and dagger carvings on sarsen surfaces, in addition to the 44 previously known [3]. The identification of further rock carvings (also called petroglyphs or engravings) at Stonehenge could provide new data to improve models of understanding of the monument. For exam-

ple, new Prehistoric carvings in addition to those identified on the exterior faces of Stones 3, 4, 5, the southwest face of 23, and the interior face of 53 (Fig. 1), could prompt further investigation into the significance of and associations between the carved stones [4]. Carvings above head height may indicate Neolithic, rather than solely Early Bronze Age, origin [5], whilst verified overlaps between carvings could help establish a chronological sequence for the rock art [6].

Abbott and Anderson-Whymark [3] in 2012 identified plane shading and luminance lensing as the most effective carving visualisation techniques for carvings at Stonehenge. The former highlights surface topography by moving a greyscale band across a point cloud representation, while the latter adjusts material luminance on a mesh model based on its position relative to a virtual lens. Each point in a point cloud represents a spatial measurement on a surface, while mesh data explicitly connects these points into a surface. Anati and Gomes [6] subsequently published re-

* Corresponding author.

E-mail address: gavinwleong@gmail.com (G. Leong).




Fig. 1. Plan of Stonehenge, with the position of upright stones with Early Bronze Age carvings indicated. Modified from Abbott and Anderson-Whymark [3]. Inset: location of Stonehenge in the UK.

sults from a 1975 survey using grazing light photography, polyvinyl sheet tracings, and latex moulds. They recorded more carvings on Stones 3, 4, 5, and 53 than Abbott and Anderson-Whymark [3], identifying 181 axe-heads, a dagger, and, for the first time, overlapping motifs. However, as discussed in section 5.3, there are limitations to the methods used by these authors.

More recent development and application of carving visualisation methods have centred on (i) trend removal, (ii) visibility-based algorithms, (iii) shaders, and (iv) mesh smoothing. Trend removal involves applying a low-pass filter to a digital elevation model to create a local elevation model, highlighting smaller-scale features [7]. The radius of the filter can be adjusted to the scale appropriate for highlighting carvings [8]. Cassen et al. [9], Pires et al. [10], and Cerrillo-Cuenca et al. [11] each created variations on trend removal by comparing a high resolution version of the engraved surface with a simplified version. Openness, a visibility-based algorithm, quantifies the degree of enclosure at each point on a surface by averaging zenith angles along profiles in eight compass directions [12]. This highlights local concavity or convexity and successfully revealed irregular carved surfaces [13]. Radiance scaling [14], a shader that correlates reflected light from a surface to the surface convexities and concavities, was recently used to visualise carvings [15]. Laplacian smoothing, a mesh manipulation technique, was used to remove general surface curvature and enhance carvings [16]. Each point is shifted towards the weighted average of its neighbours, and deviations from the original surface are computed and mapped in 2-D.

Although not designed for carving visualisation, MeshNet [17] is a supervised machine learning algorithm [18] that classifies mesh-based shapes by extracting spatial and structural features and mapping them to class labels via a neural network. At 91.9 % accuracy on the ModelNet40 benchmark [19], it may be sufficiently noise-resilient for semi-automatic carving detection. Bai et al. [20] used a UNet-based model with Sobel and Gaussian filters for carving segmentation on 2D Pitoti petroglyph images, achieving 93.5 % accuracy [21,22]. Jalandoni et al. [23] applied various models to RGB images of pictographs (painted rock art), with ResNet reaching 88 % accuracy. Horn et al. [24] trained a Faster R-CNN to detect carving motifs in trend-removed 2D images using bounding boxes. While these studies demonstrate the effective use of machine learning for 2D rock art analysis, Bai and Horn relied on clearly visible petroglyphs, and Jalandoni used pictographs, neither




Fig. 2. Photograph of the northwest face of Stone 53.

of which are directly relatable to the heavily eroded carvings at Stonehenge.

2. Research aim

The aim of this paper is to develop novel approaches to carving visualisation and recognition that improve upon the noise resilience of previous methods, so that highly faded carvings can be identified. As a case study, photogrammetric data of Stonehenge obtained in parallel to the laser scan analysis by Abbott and Anderson-Whymark [3] is used, with the northwest face of Stone 53 acting as the point of comparison. Photogrammetry data is used because it can produce higher resolution point clouds than laser scanning [25]. The three methods tested in this study are difference of Gaussians (DoG), pseudo-depth mapping (PDM), and MeshNet. This paper concludes with (i) a comparison of visualisations for the first two methods to openness [12]; radiance scaling [14]; the implementation of ‘distance between meshes’ (DBM) from Cerrillo-Cuenca et al. [11], a trend removal method; and extended difference of Gaussians (XDoG) [26], a variant of DoG; (ii) recommendations for the application of all three methods; and (iii) a discussion of the implications these methods could have on our understanding of the complexity of pattern mark-making during the Early Bronze Age at Stonehenge.

3. Materials and methods

3.1. Preparation of photogrammetric data

The northwest face of Stone 53 (Fig. 2) (decimal degrees: 51.178806, 1.826139) was chosen to test the relative effectiveness of the DoG, PDM, and MeshNet methods. It features 43 of the 45 Early Bronze Age carvings on the stone (the remaining two being located on the southwest side), ranging from a ~30 cm deeply carved dagger to many lightly carved axe-head motifs (Fig. 3). Other surface artefacts include two parallel, diagonal silica veins [27] raised in relief that cut across the northwest face of the stone as light-grey lines (Fig. 3, top), a large natural crevice in the bottom right of the surface (Fig. 2), and modern graffiti seen as Lattinate text in the upper half of the surface (Fig. 2). These natural and modern surface artefacts present challenges to all three methods as they are large features that could obstruct the enhancement or recognition of carvings in their vicinity.

The photogrammetric data for Stone 53 used in this study was captured using a Leica V-LUX 2 camera [3]. The DoG and PDM methods require the data to be processed into a 3-D mesh. This

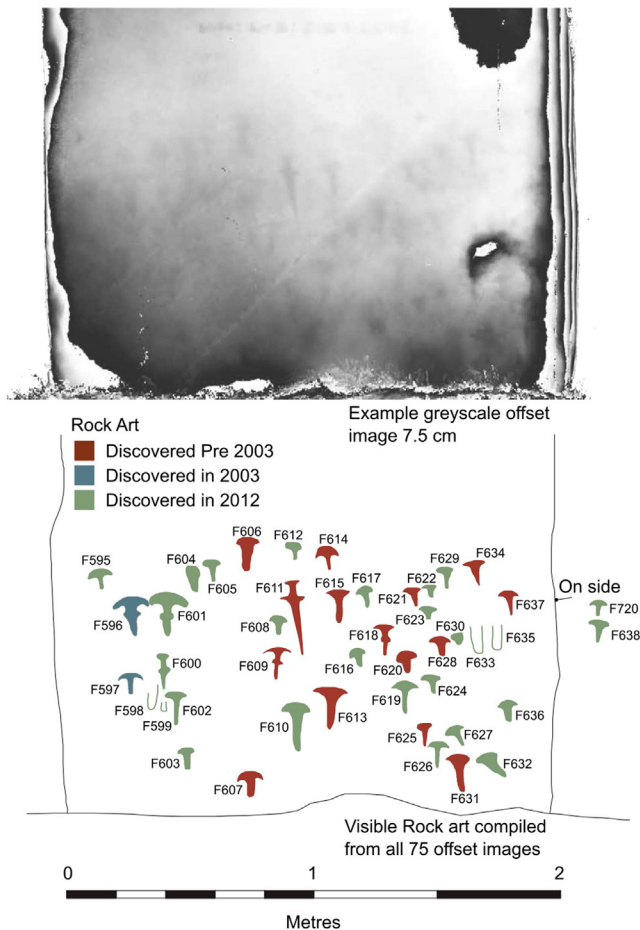


Fig. 3. Carvings on the northwest face of Stone 53 at Stonehenge. Top: greyscale plane shading visualisation of the face created by Abbott and Anderson-Whymark [3]. Bottom: their labelling of carvings visible in the visualisation. Recreated here [with permission].

was accomplished using structure from motion (SfM) photogrammetry software Agisoft Metashape [28], using the highest quality and face count settings, resulting in a mesh with 6 million vertices at 0.1 mm resolution and 1.53 mean vertices/mm² surface density. The data requirements of MeshNet differ significantly from those of DoG and PDM. In general, machine learning-based methods increase in accuracy the more data they are trained on [29]. To train MeshNet, mesh data of all carvings found at Stonehenge, processed the same way as Stone 53, were included. This meant including additional carvings from Stones 3, 4, 5, and 23 (Fig. 1). From here on, for conciseness, ‘Stone 53’ is used to refer to the ‘northwest face of Stone 53’. The following section describes the implementation of the three methods.

3.2. Application of difference of Gaussians to stone 53

The first carving visualisation method presented, DoG, removes the general curvature in a depth map by subtracting one Gaussian blurred version of the depth map from another one with less Gaussian blur. This acts to enhance small-scale features indicative of carvings.

The difference between two Gaussian filters, one with a larger standard deviation than the other, is known as DoG. DoG is more commonly used for object recognition in the scale-invariant feature transform (SIFT) computer vision algorithm [30] and for edge detection in images [31,32]. Its application in the latter makes DoG a strong candidate for carving detection in depth map images.

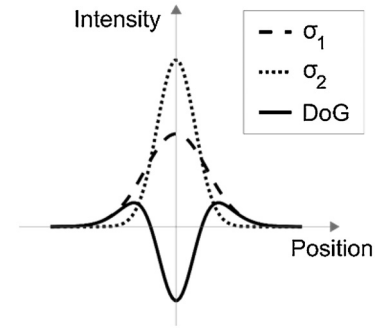


Fig. 4. 1-D representation of the difference of Gaussians (DoG) filter. A Gaussian filter with σ_2 standard deviation (dotted) is subtracted from a Gaussian filter with σ_1 standard deviation (dashed), where $\sigma_1 < \sigma_2$, to produce the DoG filter (solid).

The method of calculation for DoG is given in the below Equations [31,33]:

$$D(x, y) = \frac{1}{2\pi\sigma_1^2} e^{-\frac{x^2+y^2}{2\sigma_1^2}} - \frac{1}{2\pi\sigma_2^2} e^{-\frac{x^2+y^2}{2\sigma_2^2}}, \quad (1)$$

Where x and y describe the depth map pixel row and column coordinates and σ denotes the standard deviation, where $\sigma_1 < \sigma_2$. The term σ changes the blurring radius of the filter, the adjustment of which allows details of different scales to be revealed in the depth map. For a given pixel, its DoG value is strongly attenuated by the immediate neighbouring pixel values, due to the DoG curve falling below zero intensity, while further out its DoG value is weakly amplified as the curve sits above zero intensity (Fig. 4). To find optimal values for σ_1 and σ_2 , a grid search was used to explore the values 2.5, 5, 7.5, 10, 15, and 20 for σ_1 and 3.5, 6, 8.5, 11, 16, and 21 for σ_2 .

To obtain a depth map for Stone 53, its mesh was centred and oriented in MeshLab such that the built-in depth map shader [34] could assign a topographically descriptive greyscale colour to each face of the mesh. A high resolution 9600 × 4550 pixel image of the depth map was exported from MeshLab into MATLAB [35] to implement DoG. After empty background rows and columns of pixels were removed from the exported image, Eq. (1) was used to create a DoG filter truncated within a square convolution matrix of W^2 pixels, where $W = 2\lceil 2\sigma \rceil + 1$, with ‘ $\lceil \dots \rceil$ ’ denoting the ceiling function. This was then convolved with the imported depth map image. To ensure a high contrast resultant image, the overall greyscale image mean, μ_p , and standard deviation, σ_p , were calculated. Then, all greyscale pixels with values within $[\mu_p - \sigma_p, \mu_p + \sigma_p]$ were remapped to [0, 1]. All values below the minimum were set to 0 (black) and all above maximum set to 1 (white). From Stone 53 mesh to final visualisation, DoG was implemented in less than two hours.

3.3. Application of pseudo-depth mapping to stone 53

The second carving visualisation method, PDM, was created to address a noted limitation of DoG; the global curvature-reducing quality of DoG also reduces the contrast throughout the resultant image [36], diminishing the visual clarity of prehistoric carvings. PDM simulates the action of fitting a surface to the rock carving panel to model its global curvature. The resulting mesh is known as the *reference*. PDM then performs depth mapping of the rock carving panel, known as the *original*, relative to the reference to capture small-scale surface variations indicative of prehistoric carvings.

Subsampling is key to PDM, as it simplifies the original, removing surface features indicative of carvings while preserving the general contour. Octree subsampling [37] was applied to uni-

formly reduce and redistribute vertices, prioritising the representation of features on the scale of the entire rock surface (~10 cm to >100 cm) over those of prehistoric carvings (~0.1 cm to ~10 cm). This process removed edge and face data, requiring remeshing using Delaunay triangulation [38] followed by random point sampling to interpolate between vertices and upsample the mesh. This smoothing reduced the influence of carvings and small surface variations. Finally, depth mapping was implemented by calculating the orthogonal distance from each vertex to the nearest triangle in the reference, or to the nearest edge if the projection fell outside the triangle [39].

To implement PDM, the reference for Stone 53 was firstly created. This involved a sequence of (i) subsampling, labelled S; (ii) meshing, labelled M; and (iii) random sampling, labelled R. The Stone 53 mesh was processed in CloudCompare [40], with S executed using the 'Subsample' tool; M, the 'Delaunay 2.5D (best fitting plane)' function; and R, the 'Sample points' tool. The arrangement of the sequence of steps was 'SMRMSMRM'. The first set of 'SMRM' steps was performed with S at a finer level of octree subsampling than the S of the second set, to gradually reduce the resolution of the original. This was found in initial testing to improve the conformity of the reference to the original. Depth mapping was simulated between the reference and the original using the 'Cloud/Mesh Dist' tool [40], which calculates the orthogonal distance between a point from the original to the closest triangle in the reference. The faces of the original were then coloured in greyscale according to these distances. PDM facilitates high contrast visualisation of the original by removing saturated dark and light regions inherent to depth maps (see later Fig. 6, inset). See Appendix A for parameter values used to implement PDM. From mesh model to final visualisation, PDM was implemented in less than ten minutes.

3.4. Application of MeshNet to stone 53

The third method is a carving recognition technique that employs a neural network trained on mesh-based 3-D digital representations. Rather than enhancing the visualisation of carvings, it outputs the probability that a queried region contains a carving. This method utilises MeshNet, a preexisting 3-D shape classifier that learns shape representations [17].

MeshNet captures distinguishing features of a mesh via spatial and structural descriptors. The former uses the coordinates of each triangle centre, while the latter, each triangle corner, triangle normal, and the indices of neighbouring triangles. These descriptors incorporate kernels to model nonlinear relationships, such as the similarity between triangle corners. The features are then processed through mesh convolution blocks to aggregate neighbouring information, followed by a pooling function to estimate the likelihood that a sample belongs to a specific class, such as whether the queried region contains a carving. For a detailed explanation see Feng et al. [17]. MeshNet was trained on ModelNet40, a benchmark database of 12,311 mesh models across 40 categories [41], achieving 91.9 % accuracy.

This study trained MeshNet on mesh models of all Stonehenge carvings, as labelled by Abbott and Anderson-Whymark [3]. To enable MeshNet to differentiate between carving and non-carving regions, non-carving meshes were extracted from the same stone faces as the carvings. Approximately 90 % of the carving and non-carving meshes were used to train MeshNet, with the remaining ~10 % reserved for testing. Each mesh was simplified to 1024 faces using the 'Quadric Edge Collapse Decimation' filter from MeshLab [34] to enable batch processing. Meshes were standardised by centring, rotating, and scaling to align the centre with the origin, the average normal with the positive z-axis, and the scale with real-world dimensions. These operations were performed using a cus-

tom Blender [42] Python script. Due to the large disparity in frequency between axe-head and dagger carvings at Stonehenge (115 axe-heads vs. three daggers), and their similar forms, both types were combined in the training and testing of MeshNet.

The learning rate and weight decay hyperparameters of MeshNet were optimised to maximise carving recognition accuracy. A grid search was performed over the values 0.0005, 0.001, 0.002, 0.003 for both parameters. Performance metrics were recorded for each combination of these values. Details of the hyperparameter optimisation process are provided in Appendix B.

The success of each hyperparameter set was determined using precision, recall, F-score, and accuracy [43]. Precision measures quality by penalising misclassifications of non-carving meshes as carvings, while recall measures quantity by penalising misclassifications of carving meshes as non-carvings. F-score, the harmonic mean of precision and recall, provides a measure of balanced accuracy, whereas accuracy is the measure of all classifications that were correct, whether positive or negative, as a proportion of total classifications. These metrics were used first to identify the optimal hyperparameter set and then to assess the performance of MeshNet in labelling carvings on Stone 53.

A summary of the MeshNet implementation is provided in Appendix C. The entire process, from Stone 53 mesh creation to hyperparameter optimisation was completed in under ten hours.

4. Results

4.1. Results of difference of Gaussians on stone 53

The results of the DoG visualisation ($\sigma_1 = 5$, $\sigma_2 = 6$, and $W = 50$) are shown in Fig. 5. The method successfully identified all previously known carvings on Stone 53, matching the performance of greyscale plane shading [3]. Moreover, DoG enhanced carving clarity, leading to the discovery of two new carvings (C1 and C2), labelled 'Discovered in 2024' in Fig. 5 and confirmed via depth mapping (Fig. 6). These additions raise the total carvings on the northwest face of Stone 53 from 43 to 45. Additionally, four partial carvings (F598, F599, F633, and F635) have revised outlines, and the position of F627 has been reinterpreted.

Four further axe-head carvings (F596, F601, F609 and F610) were tentatively re-interpreted as superimposed motifs, where the smaller of each design was labelled 'Alternative interpretation' (Fig. 5). For F596, F601, and F609, lobes along the body of the axe-heads, which were previously interpreted as flanges (Fig. 3), were reinterpreted as blade-edges of superimposed axe-heads. For F610, a distinctly deeper small axe-head form (F610b) is present in the upper left-hand side of the large shallow axe-head motif (F610a). The stratigraphic relationship between the larger and smaller superimposed carvings is unclear in all four cases. Each 'Alternative interpretation' carving is typified by a darker region when visualised via depth mapping (Fig. 6). This indicates surfaces that are more deeply carved relative to the immediate surrounding surfaces, perhaps suggesting that these regions were carved into on more than one occasion.

In addition to the clear carvings visualised by DoG on Stone 53, ten other regions were highlighted as zones with potential carvings, labelled 'Potential area' in Fig. 5. In all cases the depth maps (Fig. 6, inset) revealed shallow indentations consistent with a carving, although a motif could not be clearly identified. It is possible that these areas reflect the position of very shallow or eroded carvings that require further work to define and identify.

4.2. Results of pseudo-depth mapping on stone 53

The results for the PDM visualisation are shown in Fig. 7. In addition to the two 'Discovered in 2024' carvings identified us-

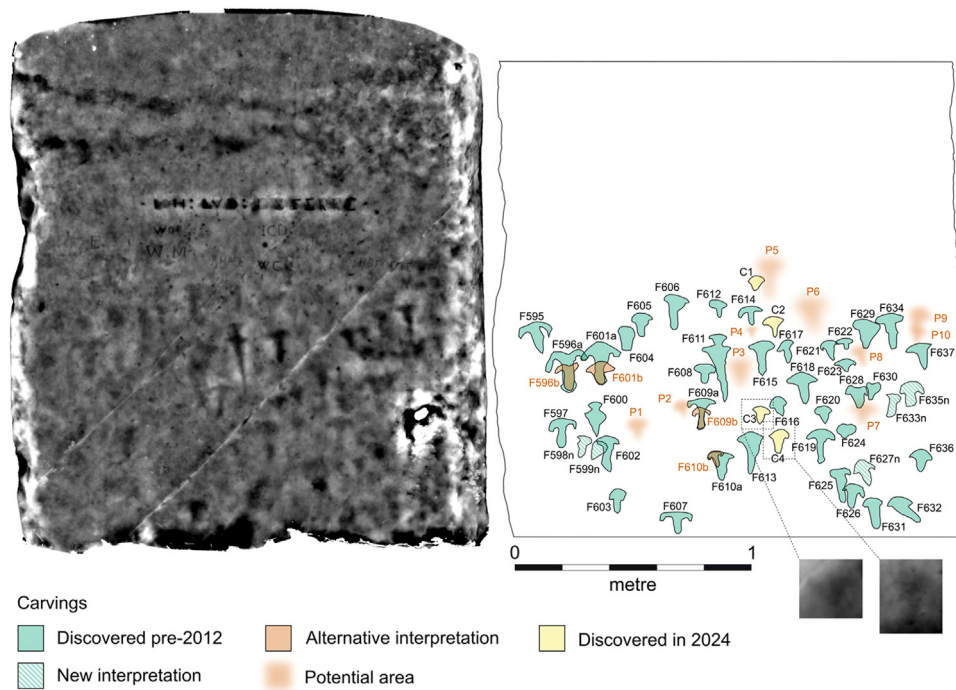


Fig. 7. Carvings identified on Stone 53 using the pseudo-depth mapping (PDM) method. Left: best visualisation of Stone 53 using PDM. Right: manually labelled carvings interpreted from the above PDM visualisation. Inset: depth map visualisation 2-times magnified.

ing DoG, two further carvings were discovered with PDM (C3 and C4) (Fig. 7, inset). These were not clearly defined in the DoG visualisation (Fig. 5, left). This raises the total carvings on Stone 53 to 47. In contrast to the DoG method, which produces as output a greyscale image, the output of the PDM method is a greyscale mesh. The colouring largely removes the global surface curvature from the depth map-like visualisation, while maintaining a high contrast image. This combines the best elements of DoG and depth mapping.

4.3. Results of MeshNet on stone 53

Fig. 8 shows the results of the MeshNet training on Stone 53, with carvings depicted using a confidence-based colour map. MeshNet correctly labelled 40 out of the 43 carvings using a hyperparameter configuration of learning rate 0.002, weight decay 0.0005, max epoch 100, and batch size 40, achieving an F-score of 91.0 % and accuracy of 90.7 %. The confusion matrix and receiver operating characteristic (ROC) curve is given in Appendix D. Performance on Stone 53 aligns closely with the benchmark dataset ModelNet40 (91.9 %) [17,41]. A recall (93.0 %) higher than precision (90.8 %) highlights the bias of MeshNet towards correctly labelling carvings over non-carvings, reducing missed carvings at the cost of more false positives.

MeshNet was trained using carvings labelled by Abbott and Anderson-Whymark [3]. While newly discovered carvings C1-C4 are highly certain, their subtlety may render them indistinguishable from non-carving meshes to MeshNet. However, this may be inconsequential, as MeshNet appears to rely on indicators different from those used by Abbott and Anderson-Whymark. Notably, confidence in labelling does not strongly correlate with carving clarity or depth. For example, F596, though deeply carved (Fig. 6, inset), is misclassified as a non-carving (Fig. 8), and F615, visible by eye (Fig. 2), receives only moderate confidence as a carving (Fig. 8).

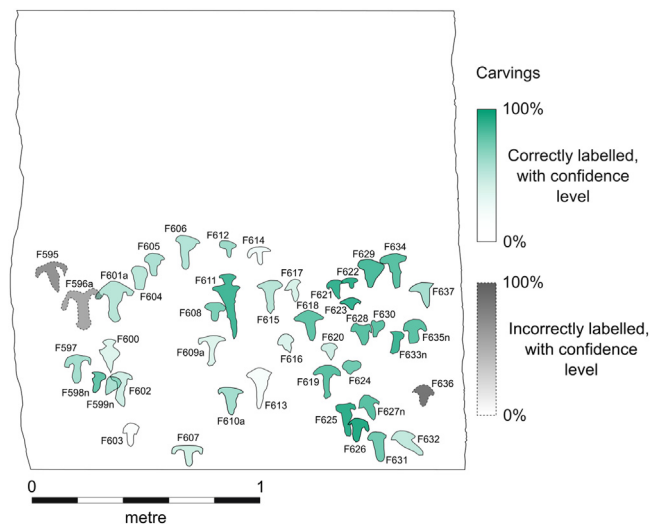


Fig. 8. MeshNet test visualisation results on Stone 53 with shade of colour indicating confidence level of the prediction by MeshNet. Carving drawings from the difference of Gaussians results is used. Three carvings were incorrectly labelled out of 43.

5. Discussion

This study implemented and compared three novel methods for carving visualisation and recognition, using the northwest face of Stone 53 at Stonehenge as a case study. The DoG method identified two previously unknown carvings, ten potential areas of carving, nine alternative interpretations of prior carvings, and all 43 carvings documented by Abbott and Anderson-Whymark [3]. PDM further revealed two additional previously unidentified carvings. MeshNet, a 3-D shape classifier, achieved 90.7 % accuracy in distinguishing carved from non-carved surface regions.




Fig. 9. Comparison of radiance scaling, openness, distance between meshes, and extended difference of Gaussians to difference of Gaussians and pseudo-depth mapping in this study. Inset for each visualisation is a 3-times magnification of a region with modern 'WC' graffiti and the carvings C1 and P5.

5.1. Comparison of visualisations

Fig. 9 compares DoG and PDM visualisations with radiance scaling [15], openness [12], and DBM [11], selected for their prior success in visualising rock carvings, and XDoG [26], a modern variant of DoG for edge enhancement in stylistic and artistic applications. For parameters used in creating each subfigure, see Appendix E. Radiance scaling adjusts reflected light intensity based on surface curvature [14], openness uses averaged profile intersections to indicate concavity [12], and DBM is related to PDM but the former instead averages two trend-removed meshes at different Poisson Screened Reconstruction depths [44] (Section 3.3). XDoG introduces a soft threshold to the base DoG for noise resilience and a variable inhibitory strength to the larger Gaussian [26].

Radiance scaling excels at visualising modern graffiti, as shown by the detailed rendering of the 'WC' in Fig. 9, inset. However, it depicts carvings like C1 and P5 poorly (Fig. 9, inset), with indistinct outlines and confounding background noise, and struggles

with faint carvings such as F596 and F601. Openness emphasises concavities and outperforms radiance scaling in visualising graffiti (Fig. 9). Despite this, it similarly fails to reveal faint carvings like F596 and F601 and produces the most blurred edges for prehistoric carvings. DBM enhances carving visibility in the centre with higher contrast than previous methods but yields low contrast in high-curvature areas, such as the panel edges, due to its averaging process. XDoG provides high contrast and reveals C1 and P5 but exhibits echo and ringing artifacts (see inset for XDoG) that reduce carving edge definition.

DoG, while producing a similar level of noise as radiance scaling and lower contrast than openness, excels at revealing faded carvings, though at the expense of sharpness, as demonstrated in the less-defined 'WC' carving. While DBM and XDoG offer higher contrast than DoG, DoG does not suffer from the low contrast in high curvature areas seen in DBM and the ringing effects of XDoG.

PDM performs worst for modern graffiti, evidenced by the lack of clarity of the 'WC' and the deeply carved Latinate graffiti in the

Table 1

Comparison of all methods showing advantages, disadvantages, and time to implement, using an Intel Core i7-7700HQ processor, at 16 GB of RAM, and a NVIDIA GeForce GTX 1060 graphics card.

Method	Advantages	Disadvantages	Time to implement
Difference of Gaussians	<ul style="list-style-type: none"> - reliably flattens a depth map to visualise variations in all but the edge of the rock carving panel - low computer processing requirement 	<ul style="list-style-type: none"> - requires grid search over parameters - produces a low-contrast visualisation 	< 2 h
Pseudo-depth mapping	<ul style="list-style-type: none"> - very fast and simple implementation - does not require grid search over parameters or parameter tuning. - method that has resulted in the highest number of carvings found on Stone 53. - low computer processing requirement 	<ul style="list-style-type: none"> - may give inaccurate visualisation due to mesh simplification for the reference mesh. - high curvature in small regions may oversaturate the visualisation, giving poor results. 	10 min
MeshNet	<ul style="list-style-type: none"> - semi-automated method: though it has high upfront time and computer processing cost, can automate the process of carving identification. 	<ul style="list-style-type: none"> - requires modern graphics card - requires cropping of data into carving and non-carving sets - requires grid search over hyperparameters - does not visualise the outline of carvings - some knowledge in applied machine learning will help optimising to dataset - accuracy of 90.7 % means some previously found carvings [3] were not recognised. 	< 10 h

upper half of Stone 53. However, PDM is the best method for visualising prehistoric carvings in a single image, particularly faded ones like C1 and P5, which are depicted with notable clarity in Fig. 9.

5.2. Recommendations for the application of difference of Gaussians, pseudo-depth mapping, and MeshNet

PDM suffers from oversaturation at the image edge (Fig. 7, left), resulting from the low face count of the reference, which sparsely samples high-curvature regions, creating large disparities between the reference and the original. Additionally, the low-resolution reference can introduce artefacts, such as the lighter area between carvings F611 and F615 (Fig. 7, left), absent in the depth map (Fig. 6, P3). Despite its speed (10 min) and lack of kernel or radius adjustments required by DoG and other methods [13], PDM may obscure details on high-curvature surfaces. It is best suited for quickly visualising highly faded carvings, such as during large-area surveys. The entire PDM process can be automated in the command line mode of CloudCompare [45] using a batch script that follows the implementation steps (Appendix A).

The main drawback of DoG is its dampening of depth information, making deeply incised carvings appear fainter than in depth maps. It also produces a white-and-black border artefact (Fig. 5, left) caused by background grey pixels in the W -sized square filter used in the calculation. While this artefact does not hinder carving visualisation, it affects cropped regions, creating uninformative 'dead space' along the borders. DoG is ideal for visualising entire carving panels without artefacts like those from PDM. Once potential carvings are identified, depth mapping can refine their outlines for greater clarity.

MeshNet serves as a feasibility study for semi-automated carving recognition. Testing on Stone 53 demonstrated accuracy comparable to benchmark datasets [17] but lower than DoG and PDM. The appeal of MeshNet lies in its potential for automation, essential for analysing large volumes of remote sensing data. Further work should focus on automating carving recognition by systematically sampling the surface with a sliding 3-D detection window [46]. To accommodate carvings of varying sizes, a 3-D variant of an image pyramid [47] can enable variable detection windows, ensuring carvings are accurately captured without overlap or omission.

PDM offers the greatest value for archaeologists seeking to identify carvings, balancing effort and gain. While DoG pro-

duces fewer artefacts, the implementation of PDM is significantly faster. Although MeshNet enables semi-automated recognition, dataset preparation remains manual and effort-intensive. Its effectiveness may also diminish with rare motifs due to insufficient training data limiting broader application. Table 1 summarises the advantages and disadvantages of DoG, PDM, and MeshNet.

5.3. Future work

The discovery of new carvings at Stonehenge using DoG and PDM supports exploring 2-D image neural networks for automated carving recognition. Examples such as ONE-PEACE [48], InternImage [49], and SOLO [50], which can segment an arbitrary number of object instances, could train on DoG or PDM images to enable automatic carving detection.

The 0.1 mm resolution 3-D dataset from this study enables quantification of engraving techniques and their variation, as demonstrated by Dubinsky and Grosman [51] using ArchCUT3-D. This could offer insights into whether the carvings were made by one or multiple individuals, the sequence of engraving events [52], and whether certain carvings received special emphasis.

Depth maps (Fig. 6, inset) suggest that some carvings may overlap, contradicting Abbott and Anderson-Whymark's [3] view that overlapping Early Bronze Age carvings at Stonehenge do not exist, while providing credence to some of the overlapping carvings on Stone 53 as suggested by Anati and Gomes [6]. However, it questions the precision with which Anati and Gomes defined individual axe-head forms and from which inferred stratigraphic relationships between carvings [6], as direct tracing and moulding are now considered subjective, irreproducible, and insufficiently detailed [7,53,54]. Grazing light also presents limitations, offering only partial visibility due to directional lighting and uneven surfaces [10]. If overlapping carvings are proven, this could imply a more complex pattern of revisited mark-making. Future research could focus on visualising subtle depth variations within or around carvings to identify more instances of carving overlap.

Abbott and Anderson-Whymark could not analyse 23 % of the stone surfaces at Stonehenge due to coverage by *Ramalina siliquosa*, a shrubby fruticose lichen. Lichen features commonly as an occlusion at other sites with possible prehistoric carvings [55]. However, as these lichen-covered surfaces were laser scanned, future work could assess the noise resilience of DoG, PDM, and MeshNet to determine their effectiveness at detecting lichen-obscured carvings.

6. Conclusions

This study introduced three methods for visualising and recognising highly faded carvings, using Stone 53 at Stonehenge as a case study. The difference of Gaussians (DoG) method detected all previously documented carvings, two new carvings, ten potential carving areas, and nine alternative interpretations of earlier findings. Pseudo-depth mapping (PDM) revealed two additional carvings beyond those identified by DoG. MeshNet, a carving recognition method, achieved 90.7 % accuracy in identifying preexisting carvings, closely matching its performance on benchmark datasets.

DoG effectively reveals highly faded carvings and is best for enhancing entire carving panels though its visualisations have low contrast. PDM offers higher contrast and is ideal for quick, effective visualisation, but struggles with high-curvature surfaces. MeshNet, while less reliable than DoG and PDM, provides a path towards automated identification for highly eroded carvings, with a measure of confidence. It is most useful when motifs recur across multiple panels or when remote sensing data exceeds the feasibility for a human to analyse. Fully automating MeshNet requires future development of a system that combines a sliding 3-D detection window and an image pyramid to systematically divide panels.

Alternative machine learning approaches could explore 2-D image segmentation algorithms, with models trained on PDM visualisations potentially outperforming the classification accuracy of MeshNet. Finally, the noise resilience of DoG, PDM, and MeshNet makes them suitable for testing on remote sensing data to reveal carvings obscured by lichen on the surfaces of Stonehenge.

Acknowledgments

This work was supported by the EPSRC Centre for Doctoral Training in Science and Engineering in Arts, Heritage and Archaeology (EP/L016036/1).

Data availability statement

The data that support the findings of this study are available from Historic England. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the corresponding author with the permission of Historic England.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.culher.2025.07.016](https://doi.org/10.1016/j.culher.2025.07.016).

References

- [1] A. Mazel, G. Nash, C. Waddington, in: *A Coming of Age*, in: art as Metaphor, Archaeopress Publishing Ltd, 2007, pp. 1–8, [doi:10.2307/j.ctvndv9dh.6](https://doi.org/10.2307/j.ctvndv9dh.6).
- [2] T. Barnett, A. Chalmers, M. Diaz-Andreu, G. Ellis, P. Longhurst, K. Sharpe, I. Trinks, 3D Laser scanning for recording and monitoring rock art erosion, *Int. Newsl. Rock Art* 41 (2005) 25–29.
- [3] M. Abbott, H. Anderson-Whymark, *Stonehenge Laser Scan: archaeological analysis report*, 2012. <https://historicengland.org.uk/research/results/reports/32-2012>.
- [4] R. Castleden, in: *The Making of Stonehenge*, 1st ed, Routledge, 1993, pp. 209–218, [doi:10.4324/9780203419465](https://doi.org/10.4324/9780203419465).
- [5] D. Field, H. Anderson-Whymark, N. Linford, M. Barber, M. Bowden, P. Linford, P. Topping, M. Abbott, P. Bryan, D. Cunliffe, C. Hardie, L. Martin, A. Payne, T. Pearson, F. Small, N. Smith, S. Soutar, H. Winton, *Analytical Surveys of Stonehenge and its environs, 2009–2013: part 2 - the stones*, *Proc. Prehist. Soc.* 81 (2015) 125–148. <https://doi.org/10.1017/ppr.2015.2>.
- [6] E. Anati, M. Varela Gomes, *Stonehenge Prehistoric Engravings. Weapons for Ancestors or for Gods?*, Instituto de Arqueologia e Paleociências, Universidade nova de Lisboa, Lisbon, 2014.
- [7] M. Diaz-Andreu, R. Hobbs, N. Rosser, K. Sharpe, I. Trinks, *Long Meg: rock art recording using 3D laser scanning*, *Past Newsl. Prehist. Soc.* (2005) 2–6.

- [8] C. Horn, D. Pitman, R. Potter, An evaluation of the visualisation and interpretive potential of applying GIS data processing techniques to 3D rock art data, *J. Archaeol. Sci. Rep.* 27 (2019) 101971, [doi:10.1016/j.jasrep.2019.101971](https://doi.org/10.1016/j.jasrep.2019.101971).
- [9] S. Cassen, L. Lescop, V. Grimaud, G. Robin, Complementarity of acquisition techniques for the documentation of Neolithic engravings: lasergrammetric and photographic recording in Gavrinis passage tomb (Brittany, France), *J. Archaeol. Sci.* 45 (2014) 126–140, [doi:10.1016/j.jas.2014.02.019](https://doi.org/10.1016/j.jas.2014.02.019).
- [10] H. Pires, F. João, L. Gonçalves-Seco, M.J. Correia Santos, O. Sousa, Morphological residual model, A tool for enhancing epigraphic readings of highly eroded surfaces, in: S. Orlandi, V. Casarosa, R. Santucci, P.M. Liuzzo (Eds.), *Proceedings of the First EAGLE International Conference on Information Technology for Epigraph and Cultural Heritage*, Sapienza Università Editrice, Rome, 2014, pp. 133–144, [doi:10.13140/RG.2.1.1323.6246](https://doi.org/10.13140/RG.2.1.1323.6246).
- [11] E. Cerrillo-Cuenca, P. Bueno-Ramírez, R. de Balbín-Behrmann, “3DMeshTracings”: a protocol for the digital recording of prehistoric art. Its application at Almendres cromlech (Évora, Portugal), *J. Archaeol. Sci. Rep.* 25 (2019) 171–183, [doi:10.1016/j.jasrep.2019.03.010](https://doi.org/10.1016/j.jasrep.2019.03.010).
- [12] R. Yokoyama, M. Shirasawa, R.J. Pike, Visualizing topography by openness: a new application of image processing to digital elevation models, *Photogramm. Eng. Remote Sens.* 68 (2002) 257–265.
- [13] F. Monna, Y. Esin, J. Magail, L. Granjon, N. Navarro, J. Wilczek, L. Saligny, S. Couette, A. Dumontet, C. Chateau, Documenting carved stones by 3D modelling – Example of Mongolian deer stones, *J. Cult. Herit.* 34 (2018) 116–128, [doi:10.1016/j.culher.2018.04.021](https://doi.org/10.1016/j.culher.2018.04.021).
- [14] R. Vergne, R. Pacanowski, P. Barla, X. Granier, C. Schlick, Radiance scaling for versatile surface enhancement, in: *Proceedings of the 13D 2010 2010 ACM SIGGRAPH Symposium Interaction 3D Graphics Games*, 1, 2010, pp. 143–150, [doi:10.1145/1730804.1730827](https://doi.org/10.1145/1730804.1730827).
- [15] S. Peña-Villasenín, M. Gil-Docampo, J. Ortiz-Sanz, Professional SfM and TLS vs a simple SfM photogrammetry for 3D modelling of rock art and radiance scaling shading in engraving detection, *J. Cult. Herit.* 37 (2019) 238–246, [doi:10.1016/j.culher.2018.10.009](https://doi.org/10.1016/j.culher.2018.10.009).
- [16] M. Zeppelzauer, G. Poier, M. Seidl, C. Reinbacher, S. Schultze, C. Breiteneder, H. Bischof, Interactive 3D segmentation of rock-art by enhanced depth maps and gradient preserving regularization, *J. Comput. Cult. Herit.* 9 (2016) 1–30, [doi:10.1145/2950062](https://doi.org/10.1145/2950062).
- [17] Y. Feng, Y. Feng, H. You, X. Zhao, Y. Gao, MeshNet: mesh neural network for 3D shape representation, *Proc. AAAI Conf. Artif. Intell.* 33 (2019) 8279–8286, [doi:10.1609/aaai.v33i01.33018279](https://doi.org/10.1609/aaai.v33i01.33018279).
- [18] S. Das, A. Dey, A. Pal, N. Roy, Applications of artificial intelligence in machine learning: review and prospect, *Int. J. Comput. Appl.* 115 (2015) 31–41, [doi:10.5120/20182-2402](https://doi.org/10.5120/20182-2402).
- [19] L. Hoang, S.-H. Lee, K.-R. Kwon, A 3D shape recognition method using hybrid deep learning network CNN-SVM, *Electronics* 9 (2020) 649, [doi:10.3390/electronics9040649](https://doi.org/10.3390/electronics9040649).
- [20] C. Bai, Y. Liu, P. Zhou, X. Wang, M. Zhou, BEGL: boundary enhancement with gaussian loss for rock-art image segmentation, *Herit. Sci.* 11 (2023) 1–10, [doi:10.1186/s40494-022-00857-5](https://doi.org/10.1186/s40494-022-00857-5).
- [21] O. Ronneberger, P. Fischer, T. Brox, U-Net: convolutional networks for biomedical image segmentation, in: N. Navab, J. Hornegger, W.M. Wells, A.F. Frangi (Eds.), *Lecture Notes on Computer Science (Including Subseries Lecture Notes on Artificial Intelligence Lecture Notes Bioinformatics)*, Springer International Publishing, 2015, pp. 234–241, [doi:10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28).
- [22] G. Poier, M. Seidl, M. Zeppelzauer, C. Reinbacher, M. Schachl, G. Bellandi, A. Marretta, H. Bischof, The 3D-Pitoti dataset, (2017) 1–7. <https://doi.org/10.1145/3095713.3095719>.
- [23] A. Jalandoni, Y. Zhang, N.A. Zaidi, On the use of Machine learning methods in rock art research with application to automatic painted rock art identification, *J. Archaeol. Sci.* 144 (2022) 105629, [doi:10.1016/j.jas.2022.105629](https://doi.org/10.1016/j.jas.2022.105629).
- [24] C. Horn, O. Ivarsson, C. Lindhé, R. Potter, A. Green, J. Ling, Artificial intelligence, 3D documentation, and rock art—approaching and reflecting on the automation of identification and classification of rock art images, *J. Archaeol. Method Theory* 29 (2022) 188–213, [doi:10.1007/s10816-021-09518-6](https://doi.org/10.1007/s10816-021-09518-6).
- [25] P. Bryan, A. Dodson, M. Abbott, Using geospatial imaging techniques to reveal and share the secrets of Stonehenge, *Int. J. Herit. Digit. Era* 3 (2014) 69–81, [doi:10.1260/2407-4970.3.1.69](https://doi.org/10.1260/2407-4970.3.1.69).
- [26] H. Winnemöller, J.E. Kyprianidis, S.C. Olsen, XDoG: an eXtended difference-of-Gaussians compendium including advanced image stylization, *Comput. Graph.* 36 (2012) 740–753, [doi:10.1016/j.cag.2012.03.004](https://doi.org/10.1016/j.cag.2012.03.004).
- [27] J.S. Uilyott, D.J. Nash, C.A. Whiteman, R.N. Mortimore, Distribution, petrology and mode of development of silcretes (sarsens and puddingstones) on the eastern South Downs, UK, *Earth Surf. Process. Landforms* 29 (2004) 1509–1539, [doi:10.1002/esp.1136](https://doi.org/10.1002/esp.1136).
- [28] Agisoft, *Agisoft Metashape User Manual version 1.8*, (2022). https://www.agisoft.com/pdf/metashape-pro_1_8_en.pdf.
- [29] M. Torii, K. Waghholikar, H. Liu, Using machine learning for concept extraction on clinical documents from multiple data sources, *J. Am. Med. Inform. Assoc.* 18 (2011) 580–587, [doi:10.1136/amiajnl-2011-000155](https://doi.org/10.1136/amiajnl-2011-000155).
- [30] D.G. Lowe, Object recognition from local scale-invariant features, *Proc. IEEE Int. Conf. Comput. Vis.* 2 (1999) 1150–1157, [doi:10.1109/iccv.1999.790410](https://doi.org/10.1109/iccv.1999.790410).
- [31] L. Assirati, N.R. da Silva, L. Berton, A. de A. Lopes, O.M. Bruno, Performing edge detection by difference of Gaussians using q-gaussian kernels performing edge detection by difference of Gaussians using q-Gaussian kernels, 2nd International Conference on Mathematical Modeling in Physical Sciences 2013, IOP Publishing, 2014, [doi:10.1088/1742-6596/490/1/012020](https://doi.org/10.1088/1742-6596/490/1/012020).

- [32] B.Y.D. Marr, E. Hildreth, Theory of edge detection, *Proc. R. Soc. Lond. Ser. B. Biol. Sci.* 207 (1980) 187–217, doi:10.1098/rspb.1980.0020.
- [33] R. Gonzalez, R. Woods, in: *Point, Line and Edge Detection*, in: *Digit. Image Process.*, 4th ed., Pearson, Harlow, 2018, pp. 701–741.
- [34] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, G. Ranzuglia, MeshLab: an open-source mesh processing tool, in: V. Scarano, R. De Chiara, U. Erra (Eds.), *Eurographics Italian Chapter Conference 2008*, The Eurographics Association, Salerno, Italy, 2008, pp. 129–136, doi:10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136.
- [35] The MathWorks Inc., MATLAB, (2023). <https://www.mathworks.com>.
- [36] F. Akram, M.A. Garcia, D. Puig, Active contours driven by difference of Gaussians, *Sci. Rep.* 7 (2017) 1–15, doi:10.1038/s41598-017-14502-w.
- [37] D. Meagher, Geometric modeling using octree encoding, *Comput. Graph. Image Process.* 19 (1982) 129–147, doi:10.1016/0146-664X(82)90104-6.
- [38] L.P. Chew, Constrained delaunay triangulations, in: *Proceedings of the 3rd Annual Symposium on Computer Geometry SCG*, 1987, 1987, pp. 215–222, doi:10.1145/41958.41981.
- [39] D. Eberly, Distance between point and triangle in 3D, *Geom. Tools* (1999) 1–4 <https://www.geometrictools.com/Documentation/DistancePoint3Triangle3.pdf>. (accessed December 4, 2024).
- [40] D. Girardeau-Montaut, CloudCompare, (2024). <https://www.danielgm.net/cc/>.
- [41] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, J. Xiao, 3D ShapeNets: a deep representation for volumetric shapes, in: *Proceedings of the IEEE Computer Society Conference on Computer Vision Pattern Recognition*, 2015, pp. 1912–1920, doi:10.1109/CVPR.2015.7298801. 07-12-June.
- [42] Blender Documentation Team, Blender 4.3 reference manual, Blender (2024). <https://docs.blender.org/manual/en/dev/>(accessed December 4, 2024).
- [43] A. Serna, B. Marcotegui, J. Hernández, Segmentation of façades from urban 3D point clouds using geometrical and morphological attribute-based operators, *ISPRS Int. J. Geo-Inf.* 5 (2016) 6, doi:10.3390/ijgi5010006.
- [44] M. Kazhdan, H. Hoppe, Screened poisson surface reconstruction, *ACM Trans. Graph.* 32 (2013) 1–13, doi:10.1145/2487228.2487237.
- [45] D. Girardeau-Montaut, Command line mode, CloudCompare.Org (2024). https://www.cloudcompare.org/doc/wiki/index.php/Command_line_mode.
- [46] S. Song, J. Xiao, Sliding shapes for 3D object detection in depth images, in: *13th European Conference Computer Vision-ECCV 2014*, Zurich, 2014, pp. 634–651, doi:10.1007/978-3-319-10599-4_41.
- [47] P.F. Felzenszwalb, R.B. Girshick, D. McAllester, D. Ramanan, Object detection with discriminatively trained part-based models, *IEEE Trans. Pattern Anal. Mach. Intell.* 32 (2010) 1627–1645, doi:10.1109/TPAMI.2009.167.
- [48] P. Wang, S. Wang, J. Lin, S. Bai, X. Zhou, J. Zhou, X. Wang, C. Zhou, ONE-PEACE: exploring one general representation model toward unlimited modalities, *ArXiv: 2305.11172* (2023) 1–30. <http://arxiv.org/abs/2305.11172>.
- [49] W. Wang, J. Dai, Z. Chen, Z. Huang, Z. Li, X. Zhu, X. Hu, T. Lu, L. Lu, H. Li, X. Wang, Y. Qiao, InternImage: exploring large-scale vision foundation models with deformable convolutions, in: *Proceedings of the IEEE/CVF Conference on Computer Vision Pattern Recognition*, 2023, pp. 14408–14419. <http://arxiv.org/abs/2211.05778>.
- [50] X. Wang, T. Kong, C. Shen, Y. Jiang, L. Li, SOLO: segmenting objects by locations, in: *Lecture Notes Computer Science (Including Subseries Lecture Notes on Artificial Intelligence on Lecture Notes Bioinformatics)*, Springer International Publishing, 2020, pp. 649–665, doi:10.1007/978-3-030-58523-5_38.
- [51] L. Dubinsky, L. Grosman, Techné of rock engravings—the Timna case study, *J. Archaeol. Method Theory* 31 (2024) 1917–1953, doi:10.1007/s10816-024-09658-5.
- [52] M. Díaz-Guardamino, L. García Sanjuán, D. Wheatley, V.R. Zamora, RTI and the study of engraved rock art: a re-examination of the Iberian south-western stelae of Setefilla and Almadén de la Plata 2 (Seville, Spain), *Digit. Appl. Archaeol. Cult. Herit.* 2 (2015) 41–54, doi:10.1016/j.daach.2015.07.002.
- [53] J. Valdez-Tullett, S. Figueiredo Persson, Digital rock art: beyond “pretty pictures”, *F1000Res* 12 (2023) 523, doi:10.12688/f1000research.127249.1.
- [54] J.L. Lerma, S. Navarro, M. Cabrelles, V. Villaverde, Terrestrial laser scanning and close range photogrammetry for 3D archaeological documentation: the upper palaeolithic Cave of Parpalló as a case study, *J. Archaeol. Sci.* 37 (2010) 499–507, doi:10.1016/j.jas.2009.10.011.
- [55] M. Díaz-Andreu, C. Brooke, M. Rainsbury, N. Rosser, The spiral that vanished: the application of non-contact recording techniques to an elusive rock art motif at Castlerigg stone circle in Cumbria, *J. Archaeol. Sci.* 33 (2006) 1580–1587, doi:10.1016/j.jas.2006.02.010.