

Evaluation of Attention Mechanisms in U-Net Architectures for Semantic Segmentation of Brazilian Rock Art Petroglyphs

Leonardi Melo^{a,b*}, Luis Gustavo^a, Dimmy Magalhães^b, Lucciani Vieira^b, and Mauro Araújo^b

^a*Computer Vision Research Center, iCEV Institute of Higher Education, Teresina, Brazil.*

^b*iCEV Institute of Higher Education, Teresina, Brazil.*

Abstract: This study presents a comparative analysis of three U-Net-based architectures for semantic segmentation of rock art petroglyphs from Brazilian archaeological sites. The investigated architectures were: (1) BEGL-UNet with Border-Enhanced Gaussian Loss function; (2) Attention-Residual BEGL-UNet, incorporating residual blocks and gated attention mechanisms; and (3) Spatial Channel Attention BEGL-UNet, which employs spatial-channel attention modules based on Convolutional Block Attention Module. All implementations employed the BEGL loss function combining binary cross-entropy with Gaussian edge enhancement. Experiments were conducted on images from the Poço da Bebidinha Archaeological Complex, Piauí, Brazil, using 5-fold cross-validation. Among the architectures, Attention-Residual BEGL-UNet achieved the best overall performance with Dice Score of 0.710, validation loss of 0.067, and highest recall of 0.854. Spatial Channel Attention BEGL-UNet obtained comparable performance with DSC of 0.707 and recall of 0.857. The baseline BEGL-UNet registered DSC of 0.690. These results demonstrate the effectiveness of attention mechanisms for archaeological heritage digital preservation, with Dice Score improvements of 2.5-2.9% over the baseline.

KEY WORDS Semantic segmentation, Rock art, Attention mechanisms, Archaeological digital heritage, Deep learning

1 Introduction

Digital preservation of archaeological sites has become increasingly relevant due to natural and anthropogenic degradation affecting rock art heritage. In northeastern Brazil, thousands of sites are registered by the Brazilian National Historic and Artistic Heritage Institute (IPHAN) [15], including Serra da Capivara National Park (Piauí) with 1,354 documented sites, with over one thousand sites containing rock art manifestations [8, 9]. The

*Correspondence to: joao_leonardi.melo@somosicev.com

Poço da Bebidinha Archaeological Complex, located in Buriti dos Montes (Piauí), contains low-relief petroglyphs engraved on paragneiss outcrops of the Canindé Unit. These petroglyphs are distributed across approximately 2 km along the Poti River margins and exhibit geometric, zoomorphic, and anthropomorphic motifs associated with the Nordeste and Agreste traditions [18, 17, 19].

Semantic segmentation of rock art presents technical challenges arising from: (1) low contrast between engravings and rock substrate; (2) differential weathering degradation; (3) presence of lichens and mineral deposits; (4) illumination variations; and (5) temporal superposition of engravings [12]. Traditional methods based on adaptive thresholding and edge detection exhibit limitations, with false positive rates exceeding 40% [24].

Convolutional neural networks such as the U-Net architecture [23] represent state-of-the-art approaches in semantic segmentation. Bai et al. [2] introduced the Border-Enhanced Gaussian Loss (BEG-L) function for rock art segmentation, combining binary cross-entropy with Gaussian edge enhancement. The framework achieved 0.865 Dice Similarity Coefficient (DSC) on the 3D-pitoti dataset [22]. Comparative studies of architectural variations incorporating attention mechanisms on Brazilian archaeological datasets remain limited.

Archaeological sites in Brazil exhibit geological, geomorphological, and stylistic characteristics that differ from European comparative collections such as the 3D-pitoti dataset [22]. The Poço da Bebidinha Complex presents mineralogical composition and erosion patterns distinct from those observed in European limestone caves. The Nordeste tradition of Brazilian rock art incorporates geometric motifs and superposition forms that differ from representations and stylistic conventions associated with European Paleolithic art.

The contributions of this work are threefold. First, it investigated deep neural networks with attention mechanisms for rock art segmentation on petroglyphs from the Poço da Bebidinha Complex. Second, it developed an annotated dataset of Brazilian petroglyphs for digital archaeological heritage documentation. Third, it provided experimental evaluation comparing different attention mechanism architectures (gated attention with residual blocks vs. CBAM-based spatial-channel attention) for rock art segmentation.

The remainder of this paper is organized as follows: Section 2 reviews related work on rock art segmentation and U-Net architectures; Section 3 describes the experimental methodology, including dataset characteristics, implemented architectures, and training protocol; Section 4 reports quantitative results and qualitative analysis; Section 5 discusses findings and practical implications; and Section 6 presents conclusions and future directions.

2 Related Work

2.1 Rock Art Segmentation

Computational methods for rock art analysis were introduced in the 1990s with image processing techniques based on morphological filters [3]. Subsequent methods included Hough transform for line detection [6], principal component analysis for dimensionality reduction [7], and k-means clustering for pattern grouping [20].

Zeppelzauer et al. [28] applied Support Vector Machines with Haar-like features for petroglyph classification, achieving 78.4% accuracy on a dataset of 1,847 images. More recent applications employ deep learning with Convolutional Neural Networks (CNNs) for classification. Jalandoni et al. [16] utilized machine learning models for automatic identification of painted rock art, achieving 89% accuracy. Defrasne et al. [5] applied digital documentation techniques for analysis of Paleolithic rock art.

2.2 U-Net Architecture and Variants

The U-Net architecture, proposed by Ronneberger et al. [23], established a reference framework for semantic segmentation through its symmetric encoder-decoder architecture with skip connections. The structure enables the preservation of high-resolution spatial information during the segmentation process, utilizing 3×3 convolutions followed by batch normalization and ReLU activation [14].

He et al. [10] introduced residual blocks that facilitated training of deeper networks through identity connections. These blocks were subsequently refined with pre-activation design [11] for improved gradient flow. Integration of residual blocks in U-Net architectures enables more stable training and improved convergence characteristics.

Oktay et al. [21] introduced the Attention U-Net, incorporating attention gates for automatic suppression of irrelevant features.

2.3 Attention Mechanisms

Attention mechanisms were proposed by Bahdanau et al. [1] for neural machine translation and extended by Vaswani et al. [25] with self-attention. In computer vision, attention enables differential weighting of spatial or channel features.

Hu et al. [13] proposed Squeeze-and-Excitation Networks (SE-Net) with channel attention via global average pooling and two fully connected layers. Wang et al. [26] presented ECA-Net (Efficient Channel Attention), without dimensional reduction, utilizing 1D convolution with adaptive kernel.

Woo et al. [27] proposed CBAM (Convolutional Block Attention Module), combining channel and spatial attention sequentially. The channel attention module utilizes dual pooling (average and max) for spatial aggregation:

$$\mathbf{M}_c(\mathbf{F}) = \sigma(\text{MLP}(\text{AvgPool}(\mathbf{F})) + \text{MLP}(\text{MaxPool}(\mathbf{F}))) \quad (1)$$

The spatial module processes channel statistics with 7×7 convolution:

$$\mathbf{M}_s(\mathbf{F}) = \sigma(\text{Conv}^{7 \times 7}([\text{AvgPool}(\mathbf{F}); \text{MaxPool}(\mathbf{F})])) \quad (2)$$

3 MATERIALS AND METHODS

3.1 Experimental Design

This comparative study was designed to evaluate three U-Net-based architectures with BEGL loss function on a Brazilian petroglyph dataset. The experimental design comprised: (1) collection and manual annotation of 82 images from the Poço da Bebidinha site; (2) $6\times$ data augmentation applied to the training images; (3) training of three architectures (BEGL-UNet baseline, Attention-Residual BEGL-UNet, Spatial-Channel Attention BEGL-UNet) with identical hyperparameters using 5-fold cross-validation; and (4) quantitative and qualitative evaluation.

3.2 Statistical Analysis

Evaluated metrics included Dice Score, Validation Loss, Precision, Recall, F1-Score, and Pixel Accuracy. Mean values across 5-fold cross-validation were computed for all architectures. Due to the limited sample size ($N = 5$ folds), statistical significance testing was not performed, and results are presented as descriptive comparisons of model performance.

3.3 Dataset

The dataset comprises 82 images from the Poço da Bebidinha archaeological site (Piauí) with segmentation masks. Images were resized and cropped to 512×512 pixels according to neural network input requirements. Preprocessing applied min-max intensity normalization to the $[0,1]$ range.

Figure 1 documents morphological variations present in the dataset, including different petroglyph typologies, conservation states, and geometric patterns.




Figure 1: Representative Dataset Samples of Brazilian Petroglyphs. Images show morphological variations including geometric patterns (a), linear engravings (b), complex overlapped figures (c), anthropomorphic representations (d), and zoomorphic figures with differential erosion (e).

3.4 Implemented Architectures

3.4.1 BEGL-UNet (Baseline)

The baseline implementation follows the original U-Net specification [23]. The encoder path consists of four hierarchical levels with double convolutional blocks following channel progression $3 \rightarrow 64 \rightarrow 128 \rightarrow 256 \rightarrow 512$. Each block contains two sequential 3×3 convolutions with batch normalization and ReLU activations. Max pooling

operations with 2×2 kernels are applied between encoder levels to progressively reduce spatial resolution and capture multi-scale features. At the bottleneck, a convolutional block at minimum resolution (32×32 pixels) expands to 1024 channels, representing the highest level of feature abstraction where complex patterns are processed. The decoder path utilizes transposed convolutions with 2×2 kernels and stride 2 for progressive upsampling, restoring spatial resolution with sequential channel reduction: $1024 \rightarrow 512 \rightarrow 256 \rightarrow 128 \rightarrow 64$. Direct concatenation between encoder and decoder features at corresponding levels preserves high-resolution spatial information along the channel dimension through skip connections. The output layer applies a final 1×1 convolution reducing to a single channel, followed by sigmoid activation: $output = \sigma(\text{Conv}_{64 \rightarrow 1}(x))$, mapping values to the $[0,1]$ range appropriate for binary segmentation.

3.4.2 Attention-Residual BEGL-UNet

The Attention-Residual BEGL-UNet architecture integrates residual blocks with attention mechanisms and BEGL loss function for rock art segmentation. The implementation employs post-activation residual blocks based on He et al. [10] with classical design (Conv-BN-ReLU-Conv-BN) followed by residual connection, formulated as:

$$\mathbf{y} = \text{ReLU}(\mathcal{F}(\mathbf{x}, \{W_i\}) + \mathbf{x}) \quad (3)$$

where $\mathcal{F}$ represents the residual mapping implemented with identity connections for different channels via 1×1 convolution. Gated attention mechanisms are applied at each decoder level for suppression of irrelevant features, defined by:

$$g_1 = \text{Conv}_{1 \times 1}(\text{BN}(g)) \quad (4)$$

$$x_1 = \text{Conv}_{1 \times 1}(\text{BN}(x)) \quad (5)$$

$$\psi = \sigma(\text{Conv}_{1 \times 1}(\text{BN}(\text{ReLU}(g_1 + x_1)))) \quad (6)$$

$$\hat{x} = x \odot \psi \quad (7)$$

where g represents the decoder gating signal, x the encoder features, and $\odot$ denotes element-wise multiplication. The overall architecture follows a symmetric encoder-decoder structure with four hierarchical encoder levels ($3 \rightarrow 64 \rightarrow 128 \rightarrow 256 \rightarrow 512$ channels) and 1024-channel bottleneck, followed by decoder with upsampling via 2×2 transposed convolutions. The BEGL loss function combines binary cross-entropy with Gaussian edge enhancement:

$$\mathcal{L}_{\text{Residual-BEGL}} = \alpha \cdot \text{MSE}(\mathcal{G}(\nabla \text{pred}), \mathcal{G}(\nabla \text{target})) + \beta \cdot \text{BCE}(\text{pred}, \text{target}) \quad (8)$$

with parameters $\alpha = 0.001$ and $\beta = 1.0$, where $\mathcal{G}$ represents Gaussian filter ($\sigma = 0.8$) and ∇ Sobel operators for edge detection.

The architecture maintains skip connections from the original U-Net, integrating residual blocks in the encoding and decoding paths, with attention gates at concatenation points to focus on segmentation-relevant features.

3.4.3 SCA-BEGL-UNet (Spatial-Channel Attention)

This architecture integrates dual spatial-channel attention based on CBAM within hybrid gated attention blocks for rock art segmentation. The channel attention module employs CBAM implementation with 16:1 reduction ratio and parallel adaptive pooling processing, formulated as:

$$\begin{aligned}\mathbf{M}_c &= \sigma(\mathbf{F}_{avg} + \mathbf{F}_{max}) \\ \mathbf{F}_{avg} &= \text{Conv}_{1 \times 1}(\text{ReLU}(\text{Conv}_{1 \times 1}(\text{AvgPool}(\mathbf{F})))) \\ \mathbf{F}_{max} &= \text{Conv}_{1 \times 1}(\text{ReLU}(\text{Conv}_{1 \times 1}(\text{MaxPool}(\mathbf{F}))))\end{aligned}\tag{9}$$

where 1×1 convolutions replace fully connected layers to preserve spatial information. The spatial attention module processes channel statistics via dual pooling followed by 7×7 convolution:

$$\mathbf{M}_s = \sigma(\text{Conv}^{7 \times 7}([\text{AvgPool}_{channel}(\mathbf{F}); \text{MaxPool}_{channel}(\mathbf{F})]))\tag{10}$$

where pooling operates along the channel dimension to create spatial attention maps. Hybrid attention gates integrate channel and spatial attention sequentially within gated attention blocks:

$$g_1 = W_g(g), \quad x_1 = W_x(x)\tag{11}$$

$$\psi_{base} = \text{ReLU}(g_1 + x_1)\tag{12}$$

$$\psi_{channel} = \text{ChannelAtt}(\psi_{base}) \odot \psi_{base}\tag{13}$$

$$\psi_{final} = \text{SpatialAtt}(\psi_{channel}) \odot \psi_{channel}\tag{14}$$

Residual blocks follow standard post-activation architecture with identity connections for different channel numbers via 1×1 convolution with batch normalization. The BEGL loss function follows standard implementation combining binary cross-entropy with Gaussian edge enhancement:

$$\mathcal{L}_{SCA-BEGL} = \alpha \cdot \text{MSE}(\mathcal{G}(\nabla pred), \mathcal{G}(\nabla target)) + \beta \cdot \text{BCE}(pred, target)\tag{15}$$

with parameters $\alpha = 0.001$ and $\beta = 1.0$, where $\mathcal{G}$ applies Gaussian filter ($\sigma = 0.8$) and ∇ Sobel operators for edge detection.

The architecture integrates residual blocks and dual attention mechanisms, which operate complementarily by applying attention over spatial and channel features in attention gates to enhance edge localization in petroglyphs.

3.5 BEGL Loss Function

The BEGL loss was implemented following the specification of Bai et al. [2], combining binary cross-entropy with Gaussian edge enhancement:

$$\mathcal{L}_{BEGL} = \alpha \mathcal{L}_G + \beta \mathcal{L}_{BCE} \quad (16)$$

where the parameters used are $\alpha = 0.001$ and $\beta = 1.0$ for all architectures. The components are defined as:

$$\mathcal{L}_{BCE} = \text{BinaryCrossEntropy}(\hat{y}, y) \quad (17)$$

$$\mathcal{L}_G = \text{MSE}(\mathcal{G}(\nabla \hat{y}), \mathcal{G}(\nabla y)) \quad (18)$$

The gradient operator ∇ utilizes Sobel filters implemented via 2D convolution with padding=1:

$$S_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, \quad S_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \quad (19)$$

$$|\nabla I| = \sqrt{(S_x * I)^2 + (S_y * I)^2 + \epsilon} \quad (20)$$

where $\epsilon = 1 \times 10^{-8}$ prevents division by zero. The Gaussian filter $\mathcal{G}$ is implemented as a separable 2D kernel with parameters kernel_size=5 and $\sigma = 0.8$:

$$g[i] = \frac{\exp(-(i - k/2)^2 / (2\sigma^2))}{\sum_j \exp(-(j - k/2)^2 / (2\sigma^2))} \quad (21)$$

$$\mathcal{G}(I) = (g \otimes g^T) * I \quad (22)$$

where k represents the kernel size and $*$ denotes 2D convolution with grouped channels for independent per-channel processing.

3.6 Experimental Protocol

3.6.1 Data Augmentation

Data augmentation techniques were implemented via the Albumentations library [4]. Geometric transformations included horizontal and vertical flips ($p=1.0$), while elastic deformations were applied with parameters $\alpha = 120$ and $\sigma = 6$ ($p=1.0$). Optical distortions utilized distortion limit=2 and shift limit=0.5 ($p=1.0$), combined with grid distortion ($p=1.0$). A $6\times$ augmentation factor was applied to the training set, whereas test set samples were maintained without modifications for validation. Figure 2 presents applied transformations, demonstrating preservation of morphological characteristics of petroglyphs.




Figure 2: Data Augmentation Examples. Representative transformations applied to training set: (a) original image, (b) horizontal flip transformation, (c) elastic deformation with parameters $\alpha=120$ and $\sigma=6$.

3.6.2 Training Configuration

The learning rate was set to $\eta = 1 \times 10^{-4}$. Adam optimizer was employed with parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Due to memory constraints, batch size was limited to $B = 2$. BEGL loss parameters were configured as $\alpha = 0.001$ and $\beta = 1.0$, with weight decay $\lambda = 1 \times 10^{-6}$.

3.6.3 Implementation Details

All models were implemented in PyTorch and trained on an NVIDIA RTX 4070 Super GPU with 12GB VRAM. Random seeds were fixed at 42 for NumPy, PyTorch, and CUDA to ensure reproducibility. Training duration was approximately 3-4 hours per architecture for 50 epochs with the specified batch size. The implementation utilized standard PyTorch libraries including torch.nn for network layers, torch.optim for optimization, and scikit-learn for cross-validation and metric computation.

4 RESULTS

4.1 Quantitative Analysis

Table 1 presents consolidated results from 5-fold cross-validation for all three architectures.

Table 1: Architecture Performance Comparison (5-fold Cross-Validation Results)

Architecture	Val Loss	DSC	Precision	Recall	F1-score	Pixel Acc
BEGL-UNet	0.077	0.690	0.617	0.834	0.690	0.968
Attention-Residual BEGL-UNet	0.067	0.710	0.629	0.854	0.710	0.967
Spatial-Channel Attention BEGL-UNet	0.081	0.707	0.618	0.857	0.707	0.968

All values represent mean performance across 5 folds at epoch 30.

4.1.1 Comparative Performance

Attention-Residual BEGL-UNet achieved the best overall performance with a DSC of 0.710 compared to the baseline DSC of 0.690, representing a 2.9% improvement. This architecture also obtained the lowest validation loss (0.067) and highest precision (0.629), recall (0.854), and F1-score (0.710). Spatial-Channel Attention BEGL-UNet achieved DSC of 0.707, representing a 2.5% improvement over baseline, with comparable recall (0.857) to the Attention-Residual architecture. All architectures achieved pixel accuracy of approximately 96.7-96.8

4.2 Qualitative Analysis

Figure 3 presents comparative segmentation results of all three architectures on petroglyph sample 15, demonstrating visual differences in edge precision and morphological preservation.




Figure 3: Comparative Segmentation Results on Sample 15. Progressive improvement in edge definition and morphological accuracy across architectures: (a) BEGL-UNet baseline, (b) Attention-Residual BEGL-UNet demonstrating superior boundary delineation, (c) Spatial Channel Attention BEGL-UNet with comparable edge precision.

4.3 Convergence Analysis

All models demonstrated convergence by epoch 30 of training. The Attention-Residual BEGL-UNet achieved the lowest validation loss (0.067) with stable performance across folds. Spatial-Channel Attention BEGL-UNet and BEGL-UNet baseline presented validation losses of 0.081 and 0.077 respectively. The consistency across 5-fold cross-validation indicates that the architectures achieved stable training despite the limited dataset size, with attention mechanisms providing consistent performance improvements.

4.4 Failure Cases Analysis

Results analysis identified specific conditions where all architectures exhibited inferior performance. These limitations affected all three evaluated architectures equally across challenging samples. Figure 4 illustrates the main failure categories observed. Main observed limitations include:

Chronological superposition: Engravings executed in multiple temporal phases created superpositions with ambiguous boundaries. Spatial-Channel Attention BEGL-UNet achieved the highest recall (0.857), indicating improved detection capability in complex scenarios.

Mineral deposition and biological colonization: Formation of iron oxide crusts and colonization by lichen organisms partially occluded rock art patterns. All architectures presented similar challenges in these cases, with modest overall DSC values reflecting the difficulty of this segmentation task.

Low radiometric contrast: Variability in mineralogical composition of the gneiss substrate and uncontrolled illumination conditions resulted in limited spectral differentiation. Attention-Residual BEGL-UNet obtained the best precision (0.629) and balanced recall (0.854), demonstrating more robust performance across challenging conditions.




Figure 4: Representative Failure Cases. Challenging conditions affecting segmentation performance across all three architectures: (a) chronological superposition with multiple temporal phases of engraving creating ambiguous boundaries, (b) mineral deposits and iron oxide crusts occluding petroglyphs, (c) low radiometric contrast between engraving and gneiss substrate under uncontrolled illumination.

5 DISCUSSION

5.1 Performance Analysis

The three analyzed architectures exhibited performance variation in the segmentation task. BEGL-UNet baseline registered DSC of 0.690 and validation loss of 0.077. Attention-Residual BEGL-UNet obtained the best overall results with DSC of 0.710 and validation loss of 0.067, achieving the highest precision (0.629) and recall (0.854) among all architectures. Spatial-Channel Attention BEGL-UNet presented comparable DSC of 0.707 with slightly higher recall (0.857) but validation loss of 0.081. All architectures achieved pixel accuracy of approximately 96.7-96.8%. These results indicate that attention mechanisms provide modest but consistent improvements over the baseline, with the Attention-Residual architecture demonstrating the most balanced performance across metrics.

5.2 Practical Implications for Archaeological Applications

The Attention-Residual BEGL-UNet model (DSC = 0.710, pixel accuracy = 96.7%) presents characteristics relevant to archaeological documentation workflows. The balanced precision (0.629) and recall (0.854) with low validation loss (0.067) indicate consistent performance across different petroglyph morphologies. While the absolute DSC values are moderate, the improvements over baseline demonstrate the potential of attention mechanisms for rock art segmentation in challenging field conditions with variable illumination and weathering patterns.

6 CONCLUSION

This study evaluated U-Net architectures with attention mechanisms for semantic segmentation of rock art on a Brazilian petroglyph dataset. Attention-Residual BEGL-UNet presented the best overall performance, with DSC of 0.710, validation loss of 0.067, and recall of 0.854, representing a 2.9% improvement over the baseline. Spatial-Channel Attention BEGL-UNet achieved comparable performance with DSC of 0.707 and recall of 0.857. Both attention-enhanced architectures surpassed the baseline BEGL-UNet (DSC: 0.690), demonstrating the effectiveness of attention mechanisms for rock art segmentation.

6.1 Limitations and Future Work

The primary limitation of this study relates to the utilization of a dataset restricted to a single archaeological site, which may limit generalization of results to other contexts. Research perspectives include evaluation of Vision Transformers for spatial dependency modeling, development of hierarchical multi-scale processing strategies, exploration of self-supervised learning, expansion of the dataset to different Brazilian archaeological sites, automatic estimation of petroglyph dimensions, directional analysis of engraving techniques, and integration with 3D reconstruction methods for archaeological documentation.

Acknowledgments

General

The authors acknowledge Dr. Gildário Lima and Prof. Lucciani Vieira for the guidance provided and for proposing this study. Institutional support from iCEV Institute of Higher Education is also recognized.

Author Contributions

L. Melo and L. Gustavo were responsible for the conception, implementation, and execution of the entire study. D. Magalhães, L. Vieira, and M. Araújo supervised the research and provided guidance throughout the project. All authors contributed equally to the preparation and revision of the manuscript.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [2] C. Bai, Y. Liu, P. Zhou, X. Wang, and M. Zhou. Begl: boundary enhancement with gaussian loss for rock-art image segmentation. *Heritage Science*, 11(1):1–10, 2023.
- [3] R. G. Bednarik. A taphonomy of palaeoart. *Antiquity*, 68(258):68–74, 1994.
- [4] A. Buslaev, V. I. Iglovikov, E. Khvedchenya, A. Parinov, M. Druzhinin, and A. A. Kalinin. Albumentations: Fast and flexible image augmentations. *Information*, 11(2):125, 2020.
- [5] C. Defrasne et al. Digital recording and analysis of paleolithic rock art: methodological developments and new insights. *Digital Applications in Archaeology and Cultural Heritage*, 20:e00165, 2021.
- [6] A. Desolneux, L. Moisan, and J. M. Morel. *From Gestalt Theory to Image Analysis: A Probabilistic Approach*. Springer Science & Business Media, New York, 2008.
- [7] C. Fritz et al. The cave art of cussac cave (france) and its contribution to understanding gravettian symbolic behavior. *Journal of Archaeological Science*, 68:1–15, 2016.
- [8] Fundação Museu do Homem Americano. Serra da capivara national park. São Raimundo Nonato: FUMDHAM. Available at: <https://www.fumdam.org.br/visite/>, 2023. Accessed: Aug. 26, 2025.
- [9] Fundação Museu do Homem Americano. Serra da capivara national park. São Raimundo Nonato: FUMDHAM. Available at: <https://www.serradacapivara.org/>, 2023. Accessed: Aug. 26, 2025.
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016.
- [12] C. Horn, O. Ivarsson, C. Lindhé, and R. Potter. Artificial intelligence, 3d documentation, and rock art—approaching and reflecting on the automation of identification and classification of rock art images. *Journal of Archaeological Method and Theory*, 29:188–213, 2021.
- [13] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7132–7141, 2018.
- [14] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *International Conference on Machine Learning*, pages 448–456, 2015.

- [15] IPHAN. Archaeological heritage. Federal Government Portal. Available at: <https://www.gov.br/iphan/pt-br/assuntos/patrimonio-material/arqueologia>, 2023. Accessed: Aug. 26, 2025.
- [16] A. Jalandoni et al. On the use of machine learning methods in rock art research with application to automatic painted rock art identification. *Journal of Archaeological Science*, 144:105629, 2022.
- [17] M. C. S. M. Lage, W. Lage, and A. L. M. Nascimento. Landscape, engravings, conservation problems: a view on the poço da bebidinha site. *Revista Mosaico – Revista de História*, 13(2):152–170, 2020.
- [18] W. Lage. The rock engravings of the bebidinha site, buriti dos montes – piauí: documentation, analysis of visual language and survey on the general state of conservation. Master’s thesis, Federal University of Piauí, Teresina, 2013.
- [19] W. Lage. *Between embroidered rocks flows a river: a Gestalt perspective to read the past*. PhD thesis, University of Coimbra, Portugal, 2018.
- [20] A. Mourão, J. Martinho, et al. Machine learning approach for exploring rock arts through the cloud infrastructure. *IEEE International Conference on Advanced Information Networking and Applications Workshops*, pages 692–697, 2016.
- [21] O. Oktay et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.
- [22] G. Poier, M. Seidl, M. Zeppelzauer, C. Reinbacher, M. Schaich, G. Bellandi, A. Marretta, and H. Bischof. The 3d-pitoti dataset: A dataset for high-resolution 3d surface segmentation. In *Proceedings of the 15th International Workshop on Content-Based Multimedia Indexing (CBMI)*, pages 1–6, 2017.
- [23] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241, 2015.
- [24] M. Seidl and C. Breiteneder. Automated petroglyph image segmentation with interactive classifier fusion. In *Proceedings of the Eighth Indian Conference on Computer Vision, Graphics and Image Processing*, pages 1–8, 2012.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [26] Q. Wang, B. Wu, P. Zhu, P. Li, W. Zuo, and Q. Hu. Eca-net: Efficient channel attention for deep convolutional neural networks. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11534–11542, 2020.

- [27] S. Woo, J. Park, J. Y. Lee, and I. S. Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.
- [28] M. Zeppelzauer et al. Interactive 3d segmentation of rock-art by enhanced depth maps and gradient preserving regularization. *Journal on Computing and Cultural Heritage*, 9(4):1–30, 2016.