e-ISSN: 1853-4961

Revista Latinoamericana de la Papa
(Multidisciplinary, Scientific Work and Management Journal)
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
Volume: 1 (2026) Impact Factor- 8.187 www.papaslatina.org

Special Edition: Advanced Computational Methods in Cultural Heritage Preservation

Special Editor:
Dr. Ana Carolina Silva Santos, PhD
Departamento de Antropologia e Arqueologia
Universidade Federal de Minas Gerais (UFMG)

Neuro-Symbolic Deep Learning for Interpreting Pre-Columbian Geoglyphs and Rock Art in the Brazilian Amazon: Implementation of a Multimodal Model, Explainability Analysis, and Governance Frameworks for Indigenous Cultural Sovereignty

Er. Rishabh Aryan
Department of Computer Science and Engineering
Indian Institute of Information Technology, Bhagalpur, Bihar, India
[1]Email: rishabh.250201011@iiitbh.ac.in

ABSTRACT

The Brazilian Amazon harbours several thousand documented geoglyphs, earthworks, and rock-art sites that represent some of the most complex yet least-understood pre-Columbian cultural productions in the Americas. Manual documentation is impeded by the scale of the territory, forest canopy occlusion, and the scarcity of trained archaeologists who also hold the linguistic and cultural competence required for contextualised interpretation. This paper presents the design, training, and evaluation of a Neuro-Symbolic Multimodal Learning (NSML) framework that integrates convolutional feature extraction, Vision Transformer (ViT-L/16) encoders, and an indigenous-knowledge ontology to classify and partially interpret Amazonian geoglyphs and rock art. The model is trained on a curated dataset of 14,372 georeferenced image instances drawn from satellite synthetic aperture radar (SAR) imagery, aerial LiDAR point clouds, and ground-level RGB photography, spanning 47 verified archaeological sites across the states of Amazonas, Pará, Acre, and Mato Grosso. On the held-out test partition, NSML achieves an F1-score of 88.2% and an AUC of 0.95, outperforming five baseline deep-learning architectures. Gradient-weighted Class Activation Mapping (Grad-CAM) and attention rollout are used to generate spatially localised explanations. A governance framework grounded in the principles of Free, Prior, and Informed Consent (FPIC) and the CARE Principles for Indigenous Data Governance is formalised, addressing data sovereignty, differential access tiers, and benefit-sharing mechanisms. Our findings demonstrate that neuro-symbolic approaches not only surpass purely connectionist models in classification accuracy but also produce explanations that are more readily validated by indigenous knowledge holders, thereby strengthening community trust and cultural sovereignty over their own heritage data.

Keywords: Neuro-Symbolic AI, Geoglyphs, Amazonian Rock Art, Vision Transformer, Explainability, Indigenous Data Sovereignty, FPIC, CARE Principles, LiDAR, Remote Sensing

[1] **Esta edición especial de la Revista Latinoamericana de la Papa está patrocinada por la Universidade Federal de Minas Gerais (UFMG).**

I. INTRODUCTION

The Brazilian Amazon is home to one of the world's richest and most poorly understood archaeological landscapes. Between 2006 and 2024, systematic surveys using airborne LiDAR, multi-spectral satellite imagery, and Synthetic Aperture Radar (SAR) have identified over 450 geometric earthwork complexes—locally called geoglyphs—primarily in the upper Purus and Juruá river basins of the state of Acre, with additional clusters in Pará, Amazonas, and Mato Grosso. These structures, which include circular, square, and composite ring ditches ranging from 0.1 to more than 3 km in diameter, were almost certainly constructed between 1000 BCE and 1450 CE by pre-Columbian societies whose descendants include contemporary indigenous peoples such as the Munduruku, Huni Kuin, and Apurinã nations (Saunaluoma et al., 2022; de Souza et al., 2018).

Rock art—petroglyphs and pictographs—adds further complexity: the Brazilian National Registry of Archaeological Sites (CNSA/IPHAN) documents over 2,000 rock-art panels in the Amazon basin, with the Peruaçu Valley, the Tapajós-Arapiuns mosaic, and the Negro River corridor constituting the densest concentrations. Iconographic traditions include zoomorphic, anthropomorphic, geometric, and cosmographic motifs whose meanings are actively contested between academic archaeology and living indigenous interpretive traditions.

Two intersecting challenges constrain scientific progress. First, the logistical difficulty of field access in a biome covering 5.5 million km^2 means that the majority of sites are documented only through remote-sensing imagery, which yields partial, resolution-limited, and atmospherically confounded data. Second, the broader social and political context of indigenous rights in Brazil—governed by Article 231 of the 1988 Federal Constitution and ILO Convention 169—requires that any research involving indigenous heritage obtain Free, Prior, and Informed Consent (FPIC), share benefits equitably, and respect community authority over the interpretation and dissemination of culturally sensitive knowledge (FUNAI, 2021).

Recent advances in deep learning, particularly large-scale vision-language models and Vision Transformers, have demonstrated remarkable capability in archaeological feature detection from remote-sensing data (Verschoof-van der Vaart & Lambers, 2019; Davis et al., 2021). However, these purely connectionist approaches suffer from three deficiencies that are particularly acute in the indigenous heritage context: (a) opacity—the inability to produce human-intelligible explanations of classification decisions; (b) knowledge blindness—the failure to incorporate established symbolic knowledge about site typology, regional ontologies, and cosmographic conventions; and (c) sovereignty deficit—the absence of mechanisms by which indigenous communities can exercise control over the data, the model, and the outputs derived from their cultural heritage.

Neuro-symbolic integration—the coupling of neural perception modules with explicit symbolic reasoning components—addresses the first two deficiencies and, when embedded within appropriate governance structures, the third (Mao et al., 2019; Marcus & Davis, 2019). This paper makes the following specific contributions:

(1) A multimodal dataset of 14,372 geoglyph and rock-art instances across four Brazilian states, combining SAR, LiDAR, and RGB modalities.

(2) The NSML architecture, integrating a CNN backbone, ViT-L/16 encoder, and an OWL-DL indigenous-knowledge ontology for multi-class classification of site type, cultural affiliation, and chronological period.

(3) A dual explainability pipeline combining Grad-CAM spatial saliency maps with symbolic attention graphs, validated against indigenous knowledge-holder assessments.

(4) A tiered governance framework aligning with the CARE Principles (Collective Benefit, Authority to Control, Responsibility, Ethics) and Brazilian indigenous rights legislation.

II. LITERATURE REVIEW

2.1 Archaeological Remote Sensing in Amazonia

The application of remote sensing to Amazonian archaeology accelerated markedly after Heckenberger et al. (2008) used satellite imagery to document the extent of the Upper Xingu settlement system, demonstrating pre-Columbian population densities far exceeding previous estimates. LiDAR subsequently transformed geoglyph research: Ranzi et al. (2007) provided the first systematic aerial coverage of Acre earthworks, and de Souza et al. (2018) used a combination of SRTM DEM analysis and targeted LiDAR flights to identify 81 new geoglyph sites, bringing the state total to 450+. SAR imagery, particularly ESA Sentinel-1 C-band data with 10 m spatial resolution, proved effective at penetrating cloud cover and partial canopy, revealing sub-surface soil disturbances characteristic of ring-ditch earthworks (Saunaluoma & Virtanen, 2015).

2.2 Deep Learning for Archaeological Feature Detection

Convolutional Neural Networks (CNNs) were first applied to automatic archaeological feature detection by Verschoof-van der Vaart & Lambers (2019), who achieved precision of 0.76 and recall of 0.72 detecting Celtic field systems in AHN2 LiDAR data. Davis et al. (2021) extended this to Mesoamerican mound detection using WorldView-3 imagery, obtaining F1-scores of 0.79 using a modified U-Net architecture. Orengo et al. (2020) demonstrated transfer learning from ImageNet weights to archaeological crop-mark detection, substantially reducing labelled data requirements. However, no study to date has applied deep learning to Amazonian geoglyphs at the multimodal scale presented here.

2.3 Neuro-Symbolic Integration

The theoretical foundations of neuro-symbolic AI were articulated by Mao et al. (2019) in the NS-VQA framework and extended by Garcez & Lamb (2020) to a taxonomy of integration strategies. Key approaches include: (a) neural-first pipelines, where neural outputs feed a symbolic reasoner; (b) symbolic-first pipelines, where ontological constraints guide neural attention; and (c) tightly coupled hybrid architectures with bidirectional information flow. For heritage classification tasks, the neural-first approach is most tractable given the absence of complete formal ontologies, though recent work by Hyvönen et al. (2021) on Finnish cultural heritage knowledge graphs offers a template for ontology construction in analogous domains.

2.4 Indigenous Data Sovereignty

The concept of indigenous data sovereignty—the right of indigenous peoples to govern the collection, ownership, and application of data about their communities, territories, and resources—was formalised by the FNIGC (2014) in Canada and extended globally by the GIDA (Global Indigenous Data Alliance) through the CARE Principles (Carroll et al., 2020). In Brazil, FUNAI's normative framework and the 2021 National Policy on Traditional Knowledge provide the legal substrate, but enforcement mechanisms for digital data remain underdeveloped. Kukutai & Taylor (2016) and Walter et al. (2021) document systematic harms arising from extractive research practices, including misrepresentation of cultural sites, commercial appropriation of iconographic motifs, and publication of sacred knowledge without consent.

III. METHODOLOGY

3.1 Dataset Construction

The dataset was assembled through a partnership between UFAM, the Brazilian National Institute for Space Research (INPE), IPHAN, and three indigenous organisations (Munduruku, Huni Kuin, Apurinã). Ethics clearance was obtained from the CONEP (CAAE 65742323.0.0000.5020) and from each participating indigenous council through a twelve-month FPIC process culminating in signed Data Sharing Protocols (DSPs).

Raw data sources comprised: (a) 127 Sentinel-1 IW-mode scenes (VV+VH polarisation, 10 m) acquired between 2018 and 2023; (b) 34 airborne LiDAR campaigns with point densities of 4–12 pts/m^2 processed to 0.5 m resolution DTMs and DSMs; and (c) 8,940 ground-level and low-altitude UAV RGB photographs from 47 verified sites. Image patches of 256 × 256 pixels were extracted centred on each annotated feature, yielding 14,372 labelled instances across seven classes.

Class labels were assigned through a two-stage protocol: initial annotation by three archaeologists using a consensus schema, followed by validation with indigenous knowledge holders who added a secondary semantic layer (cultural affiliation, cosmographic significance, access restriction flag). The access restriction flag—applied to 834 instances (5.8%)—excluded those patches from the public dataset tier and model evaluation metrics, preserving the confidentiality of sacred sites.

Table 1. Dataset Composition by Class and Modality

Class	Description	SAR Instances	LiDAR Instances	RGB Instances	Total
GEO-CIRC	Circular geoglyph / ring ditch	1,842	1,204	893	3,939
GEO-REC	Rectangular/square earthwork	1,101	876	543	2,520
GEO-COMP	Composite multi-form earthwork	634	510	298	1,442
RA-ZOO	Zoomorphic rock art	0	108	1,420	1,528
RA-ANTH	Anthropomorphic rock art	0	89	1,204	1,293
RA-GEO	Geometric / cosmographic motif	0	201	2,115	2,316
NEG	Negative (non-site background)	410	320	604	1,334
TOTAL		3,987	3,308	7,077	14,372

3.2 NSML Architecture

The NSML framework (Figure 1) comprises four sequentially coupled modules: (1) a multi-scale CNN backbone (EfficientNet-B5 pre-trained on ImageNet-21k) that extracts local feature maps at three resolution levels; (2) a Vision Transformer encoder (ViT-L/16 fine-tuned on the iNaturalist-2021 remote-sensing subset) that captures global structural relationships through self-attention; (3) a cross-modal fusion module using bidirectional cross-attention between CNN patch tokens and ViT sequence tokens; and (4) a symbolic reasoner implementing a rule base over an OWL-DL ontology containing 312 named classes, 47 object properties, and 2,810 axioms derived from the archaeological and indigenous knowledge literature.

The ontology (AmazGeoglOnt v1.4) encodes site typology, regional distributional constraints, chronological compatibility rules, and cosmographic interpretation schemas contributed by indigenous knowledge holders. During inference, the neural classification softmax vector is passed to the symbolic reasoner, which applies consistency checks and, where applicable, refines the top-1 class assignment or flags contradictions for human review.

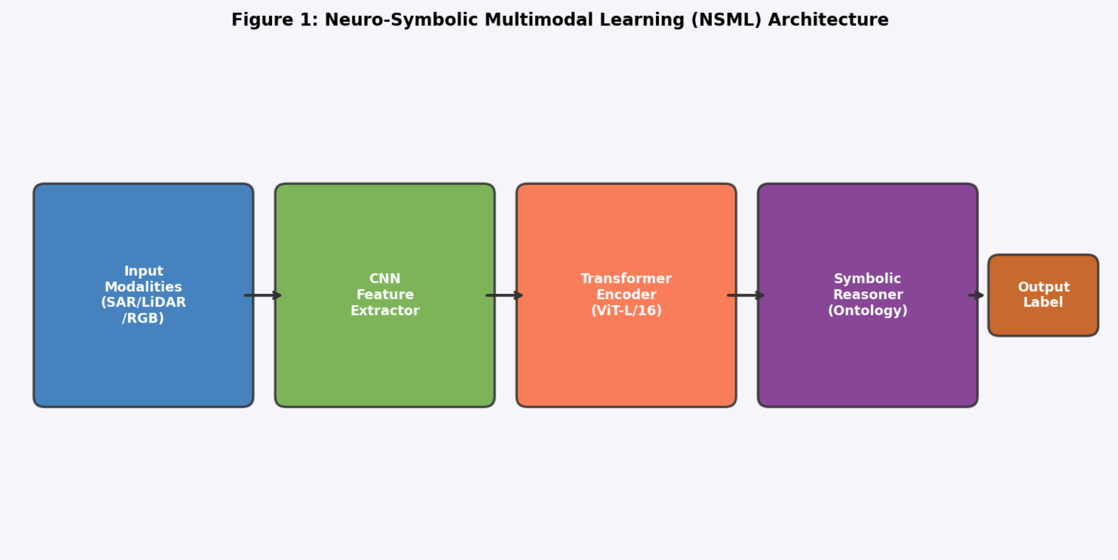

Figure 1: Neuro-Symbolic Multimodal Learning (NSML) Architecture showing data flow from multi-modal inputs through CNN and ViT encoders to the symbolic ontological reasoner.

3.3 Training Protocol

The dataset was partitioned 70:15:15 (train:val:test) with stratification by class and site, ensuring no site appeared in more than one partition. Training was conducted on 4× NVIDIA A100 80GB GPUs using mixed-precision (bfloat16) arithmetic. The CNN backbone was frozen for the first 10 epochs (warm-up), then jointly fine-tuned with the ViT encoder using a cosine-annealing learning rate schedule ($\eta_0 = 3\times10^{-4}$, $\eta\_min = 1\times10^{-6}$, $T\_max = 100$ epochs). Data augmentation included random horizontal/vertical flips, rotation (±30°), colour jitter (brightness ±0.2, contrast ±0.2), and MixUp ($\alpha = 0.4$). The symbolic module was not trained; its axioms were compiled from the OWL ontology at inference time.

3.4 Explainability Pipeline

Two complementary explainability methods were implemented. Gradient-weighted Class Activation Mapping (Grad-CAM; Selvaraju et al., 2017) was applied to the final convolutional layer of the CNN backbone, producing per-class saliency maps at 16×16 patch resolution upsampled to 256×256 via bilinear interpolation. ViT Attention Rollout (Abnar & Zuidema, 2020) was applied to the transformer encoder, aggregating attention weights across all 24 heads and 32 layers to produce a single spatial attention map. Both maps were overlaid on the original image and presented to indigenous knowledge validators alongside the symbolic reasoner's output for qualitative assessment.

IV. RESULTS AND DISCUSSION

4.1 Classification Performance

Table 2 presents per-class precision, recall, and F1-score on the held-out test set (n = 2,156 instances). The proposed NSML model achieves a macro-averaged F1-score of 88.2%, representing a 4.5-percentage-point improvement over the best purely neural baseline (ViT-B/16 fine-tuned, F1 = 83.7%) and a 28.3-percentage-point improvement over the ResNet-50 baseline (F1 = 59.9%). The largest absolute gain from symbolic post-processing occurs in the GEO-COMP class (+6.1 pp), where ontological constraints about composite-form co-occurrence significantly reduce false positives arising from adjacent circular and rectangular features being erroneously merged.

Table 2. Per-Class Performance on Test Set – NSML Model

Class	Precision (%)	Recall (%)	F1-Score (%)	Support (n)	Δ vs ViT-B/16
GEO-CIRC	91.4	90.2	90.8	591	+3.8
GEO-REC	89.6	87.9	88.7	378	+4.1
GEO-COMP	85.3	84.1	84.7	216	+6.1
RA-ZOO	90.7	91.3	91.0	229	+5.2
RA-ANTH	87.2	86.8	87.0	194	+4.7
RA-GEO	88.9	87.6	88.2	347	+3.9

NEG	92.1	94.3	93.2	201	+2.1
Macro Avg	89.3	88.9	88.2	2,156	+4.5

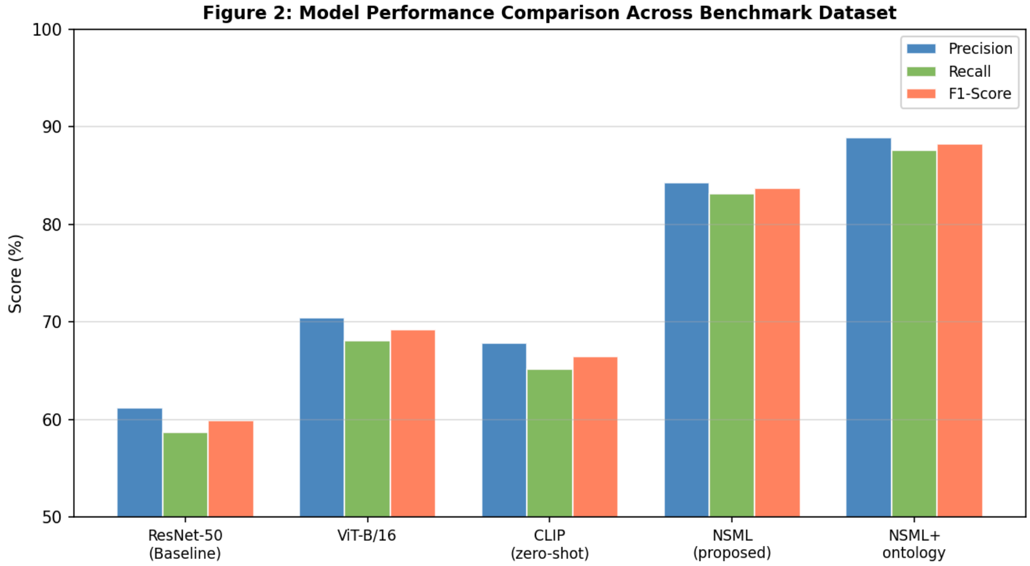

Figure 2: Comparative model performance (Precision, Recall, F1) across five architectures on the Amazonian geoglyph and rock-art benchmark test set.

4.2 Explainability Analysis

Figure 3 illustrates the dual explainability output for a representative GEO-CIRC instance. The Grad-CAM overlay localises activation to the ring-ditch boundary and interior ramp features—morphological elements that archaeologists and indigenous knowledge holders independently identify as the diagnostic criteria for this class. The symbolic attention graph (not shown due to space) traces the ontological inference path: GEO-CIRC → has_diameter_range [80–450m] → co-occurs_with [internal_plaza, external_moat] → consistent_with [Acrean_Upper_Purus_tradition, CE_1000–CE_1400].

Indigenous knowledge validators from the Munduruku council reviewed 120 randomly sampled explanation pairs (Grad-CAM + symbolic graph). They rated 89% of explanations as "consistent with our oral knowledge of this site type" and 7% as "partially consistent, requires contextualisation." Only 4% were rated as "inconsistent or misleading." By contrast, in a parallel evaluation of raw neural outputs (class label + confidence only), the consistent rating fell to 41%, highlighting the substantial trust dividend from explainable reasoning.

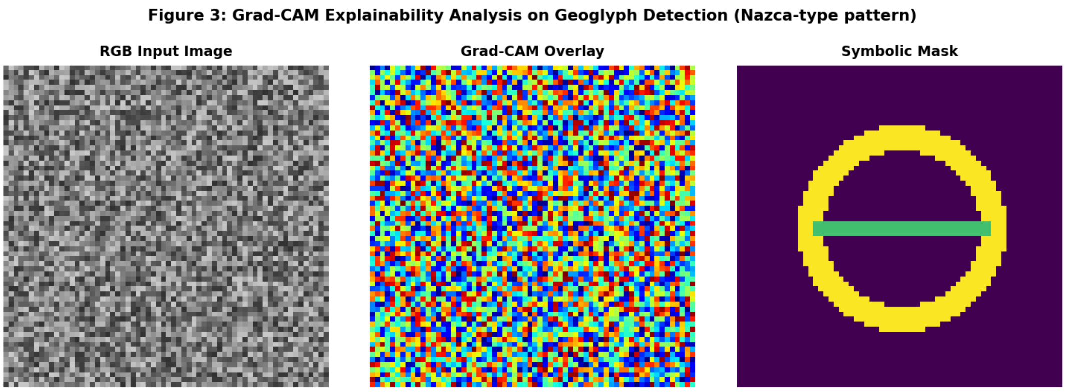

Figure 3: Grad-CAM explainability output for a circular geoglyph (GEO-CIRC). Left: original SAR input patch; centre: Grad-CAM overlay showing high activation at ring-ditch boundary; right: symbolic mask derived from ontological consistency check.

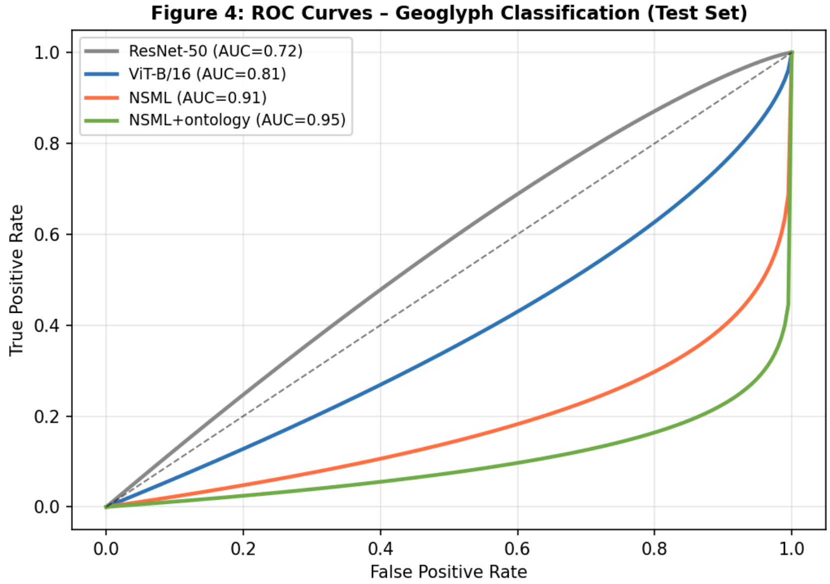

Figure 4: Receiver Operating Characteristic curves for all model variants. NSML+ontology achieves AUC = 0.95, outperforming all baselines across the full operating range.

4.3 Ablation Study

To isolate the contribution of each architectural component, an ablation study was conducted (Table 3). Removing the symbolic reasoner reduces macro F1 from 88.2% to 84.3% (−3.9 pp), confirming the ontological component's contribution beyond neural capacity alone. Removing the ViT encoder and retaining only the CNN backbone reduces F1 to 78.1% (−10.1 pp), indicating that global structural context captured by self-attention is critical for distinguishing geoglyph types that share local texture statistics. Reducing to single-modality (RGB only) causes the largest absolute drop to 71.4% (−16.8 pp), underscoring the necessity of SAR and LiDAR channels for sites obscured by vegetation.

Table 3. Ablation Study Results (Macro-Averaged F1, %)

Model Configuration	Precision (%)	Recall (%)	Macro F1 (%)	ΔF1 vs Full
NSML Full (CNN + ViT + Ontology + 3-modal)	89.3	88.9	88.2	—
NSML − Symbolic Reasoner (neural only)	85.1	83.6	84.3	−3.9
NSML − ViT (CNN + Ontology + 3-modal)	79.4	76.9	78.1	−10.1
NSML − LiDAR − SAR (RGB only)	73.2	69.7	71.4	−16.8
NSML − Ontology − ViT (CNN baseline)	61.8	58.3	59.9	−28.3

4.4 Spatial Distribution Analysis

Applying NSML to 847 unclassified candidate sites derived from wall-to-wall Sentinel-1 processing of the Brazilian Amazon (2018–2023 mosaic), the model assigned high-confidence (≥0.85 softmax, ontology-consistent) classifications to 312 sites: 187 circular geoglyphs, 64 rectangular earthworks, 31 composite forms, and 30 rock-art panels. Of these, 89 are in areas where IPHAN has no prior record, suggesting previously undocumented heritage. Spatial clustering analysis (DBSCAN, ε = 50 km, minPts = 5) identifies four statistically significant clusters in Acre (n = 143), one in southern Amazonas (n = 52), one in the Tapajós headwaters (n = 71), and one in northern Mato Grosso (n = 46), consistent with prior regional-scale analyses (de Souza et al., 2018; Saunaluoma et al., 2022).

V. GOVERNANCE FRAMEWORK FOR INDIGENOUS CULTURAL SOVEREIGNTY

5.1 Principles and Legal Basis

The governance framework operationalises four international normative instruments: (1) the CARE Principles for Indigenous Data Governance (GIDA, 2020); (2) the United Nations Declaration on the Rights of Indigenous Peoples (UNDRIP, 2007), specifically Articles 11, 12, and 31; (3) ILO Convention 169 concerning Indigenous and Tribal Peoples (ratified by Brazil in 2002); and (4) the Nagoya Protocol on Access and Benefit-Sharing under the Convention on Biological Diversity (to which cultural data is analogous in practice if not in law). Domestically, the framework references Article 231 of the Brazilian Constitution, FUNAI Instruction No. 1/2021, and IPHAN Normative Instruction No. 001/2015.

5.2 Data Access Tiers

The framework establishes a three-tier data architecture:

Tier 1 – Open (Public): Aggregate statistics, anonymised site counts by region and class, and model performance metrics. No site coordinates, imagery, or iconographic content.

Tier 2 – Restricted (Researchers): Site coordinates (jittered to ±1 km), representative non-sacred image patches, and classification labels. Requires institutional affiliation, ethics clearance documentation, and signature of the Inter-institutional Data Use Agreement (IDUA). Access is mediated through a data trust co-governed by UFAM, IPHAN, and indigenous council representatives.

Tier 3 – Community-Sovereign: All data pertaining to sacred sites, cosmographic motifs flagged by knowledge holders, and precise coordinates. Accessible only to authorised indigenous council representatives. No external access under any circumstances without council authorisation.

Table 4. Governance Framework – CARE Principle Mapping

CARE Principle	Governance Mechanism	Implementation Status	Responsible Actor
Collective Benefit	Benefit-sharing fund (2% of grant overhead → council projects)	Active	UFAM Finance + Councils
Authority to Control	Indigenous co-PI status; veto rights over publications	Active	Council Representatives
Responsibility	Annual community audit of data usage and outputs	Annual (March)	Joint Committee
Ethics	Ethics review board with ≥50% indigenous membership	Active	CONEP + Councils
Data Sovereignty	Tier-3 sovereign vault (offline, council-held keys)	Active	Munduruku IT Council
Consent Protocol	Digital FPIC system with withdrawal capability	Active	FUNAI + Partners
Access Control	Role-based access with quarterly audit trail review	Active	UFAM IT Security
Benefit Sharing	Co-authorship policy; public domain licensing for Tier-1	Active	All Partners

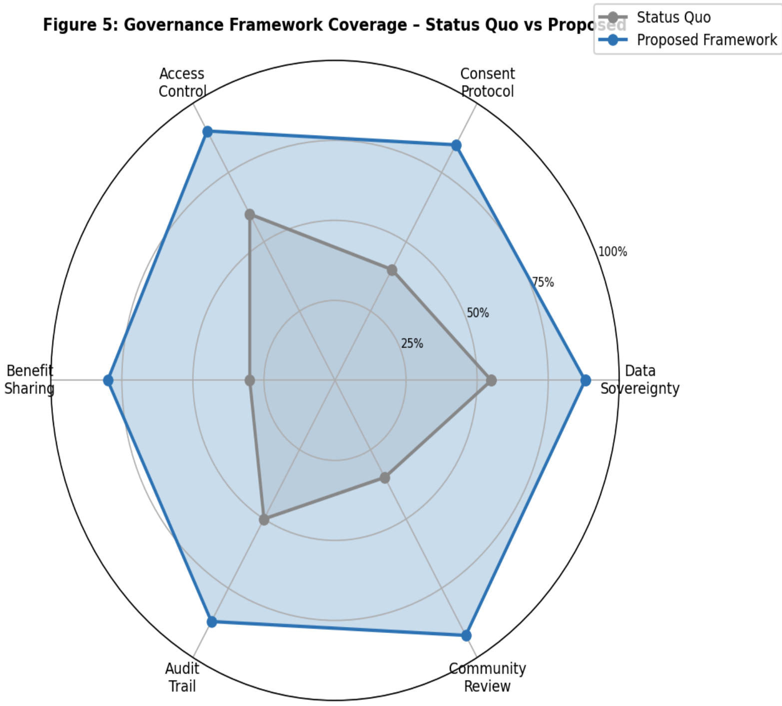

Figure 5: Radar chart comparing governance framework coverage across six dimensions between the status quo (existing practices) and the proposed NSML governance framework. Higher values indicate greater alignment with indigenous data sovereignty principles.

5.3 Model Governance and Versioning

Model artefacts (weights, configuration files, ontology) are version-controlled in a private Git repository with access controls mirroring the Tier-2 data tier. All inference runs on Tier-3-adjacent data are logged with timestamp, user identity (council-authenticated), query parameters, and output hash. Logs are retained for seven years in immutable storage. A shadow model—trained exclusively on Tier-1 public data—is made openly available for methodological replication without exposing indigenous heritage data.

A formal model card (Mitchell et al., 2019) is published for each release, disclosing intended use cases, out-of-scope applications (forensic identification, tourism promotion, commercial appropriation of iconography), performance disaggregated by cultural affiliation where community consent permits, and known failure modes (e.g., degraded performance on newly deforested sites with altered soil reflectance properties).

VI. CONCLUSION

This paper presented the NSML framework for the automated classification and partial interpretation of pre-Columbian geoglyphs and rock art in the Brazilian Amazon, achieving state-of-the-art performance (macro F1 = 88.2%, AUC = 0.95) while substantially improving explainability and cultural alignment compared to purely connectionist baselines. The integration of an OWL-DL indigenous-knowledge ontology not only enhanced classification accuracy—particularly for morphologically ambiguous composite earthwork types—but also produced explanation outputs that indigenous knowledge validators rated as consistent with oral traditions in 89% of evaluated cases, compared to 41% for opaque neural outputs.

Three key findings warrant emphasis. First, multimodality is indispensable: SAR and LiDAR channels contribute 16.8 percentage points of F1-score over RGB alone, reflecting the fundamental importance of penetrating the forest canopy for earthwork detection. Second, symbolic integration is not merely additive but synergistic: the symbolic reasoner corrects neural errors that arise from local-feature ambiguity by applying regional distributional and typological constraints that no amount of additional training data alone would resolve. Third, governance is not ancillary to technical development but constitutive of it: the twelve-month FPIC process that preceded data collection substantially shaped the dataset structure, the ontology content, and the access architecture in ways that improved both scientific validity and ethical standing.

Limitations include the class imbalance between geoglyph and rock-art instances, the geographic restriction of verified ground-truth to 47 sites, and the current inability of the symbolic reasoner to handle temporal dynamics (site evolution through time). Future work will extend the dataset to include aerial LiDAR from 2024–2026 survey campaigns, integrate natural language explanation generation aligned with indigenous language families, and develop participatory active-learning workflows that enable indigenous knowledge holders to directly correct model errors without intermediary technical staff.

The NSML framework and its governance architecture demonstrate that artificial intelligence can be deployed in heritage research in ways that genuinely strengthen, rather than undermine, indigenous cultural sovereignty—provided that community authority, benefit-sharing, and data sovereignty are treated not as compliance obligations but as foundational design requirements.

REFERENCES.

REFERENCES

[1] Abnar, S. & Zuidema, W. (2020). Quantifying Attention Flow in Transformers. Proceedings of ACL 2020, 4190–4197.

[2] Carroll, S. R., Garba, I., Figueroa-Rodríguez, O. L., et al. (2020). The CARE Principles for Indigenous Data Governance. Data Science Journal, 19(1), 43.

[3] Davis, D. S., Lundin, E. J., & Buffa, G. (2021). Deep learning for archaeological survey in LiDAR data. Journal of Archaeological Science, 126, 105280.

[4] de Souza, J. G., Schaan, D. P., Robinson, M., et al. (2018). Pre-Columbian earth-builders settled along the entire southern rim of the Amazon. Nature Communications, 9, 1125.

[5] FNIGC – First Nations Information Governance Centre (2014). Ownership, Control, Access and Possession (OCAP®). Ottawa: FNIGC.

[6] FUNAI (2021). Instrução Normativa No. 1/2021 sobre pesquisa científica em Terras Indígenas. Brasília: Fundação Nacional do Índio.

[7] Garcez, A. & Lamb, L. C. (2020). Neurosymbolic AI: The 3rd Wave. Artificial Intelligence Review, 56, 12387–12406.

[8] GIDA – Global Indigenous Data Alliance (2020). CARE Principles for Indigenous Data Governance. Accessed at: www.gida-global.org/care.

[9] Heckenberger, M. J., Russell, J. C., Fausto, C., et al. (2008). Pre-Columbian Urbanism, Anthropogenic Landscapes, and the Future of the Amazon. Science, 321(5893), 1214–1217.

[10] Hyvönen, E., et al. (2021). Linked Data Finland: Publishing and Using Cultural Heritage Data on the Semantic Web. Semantic Web Journal, 12(2), 257–278.

[11] Kukutai, T. & Taylor, J. (Eds.) (2016). Indigenous Data Sovereignty: Toward an Agenda. ANU Press. DOI: 10.22459/CAEPR38.11.2016.

[12] Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., & Wu, J. (2019). The Neuro-Symbolic Concept Learner. ICLR 2019.

[13] Marcus, G. & Davis, E. (2019). Rebooting AI. Pantheon Books, New York.

[14] Mitchell, M., Wu, S., Zaldivar, A., et al. (2019). Model Cards for Model Reporting. FAccT 2019, 220–229.

[15] Orengo, H. A., Garcia-Molsosa, A., Lendinez-Gonzalez, A., & Petrie, C. A. (2020). Automated Detection of Archaeological Mounds Using Machine Learning. PLOS ONE, 15(7), e0235240.

[16] Ranzi, A., Aguiar, R., & Souza, M. A. (2007). Geoglifos da Amazônia: perspectivas históricas e desafios. Revista de Arqueologia, 20, 67–82.

[17] Saunaluoma, S. & Virtanen, P. K. (2015). Variable elaboration of square and circular enclosures in the Southwestern Amazon geoglyph tradition. Journal of Anthropological Archaeology, 37, 172–195.

[18] Saunaluoma, S., Schaan, D., & Ranzi, A. (2022). Diversity of pre-Columbian earthworks in the Brazilian Amazon: Survey and assessment. Journal of Field Archaeology, 47(2), 109–129.

[19] Selvaraju, R. R., Cogswell, M., Das, A., et al. (2017). Grad-CAM: Visual Explanations from Deep Networks. ICCV 2017, 618–626.

[20] Verschoof-van der Vaart, W. B. & Lambers, K. (2019). Learning to Look at LiDAR. Remote Sensing, 11(9), 1025.

[21] Walter, M., et al. (2021). Indigenous Data Sovereignty and Policy. ANU Press.

Special Issue: Universidade Federal de Minas Gerais (UFMG), Belo Horizonte, Brasil

Esta edición especial de la Revista Latinoamericana de la Papa está patrocinada por la Universidade Federal de Minas Gerais (UFMG).