Towards SAR-Based Monitoring of Illegal Mining in the Brazilian Amazon Using Convolutional Neural Networks

Nelson Lemes Neto[1], Maria de Lourdes Bueno Trindade Galo[1,2], Renan Américo de Oliveira[1], Fernanda Sayuri Yoshino Watanabe[1,2], Mauricio Galo[1,2]

[1] Graduate Program in Cartographic Sciences, São Paulo State University (UNESP), Presidente Prudente, São Paulo, Brazil – {nelson.lemes, renan.americo}@unesp.br
[2] Department of Cartography, Faculty of Science and Technology (FCT), São Paulo State University (UNESP), Presidente Prudente, São Paulo, Brazil – {trindade.galo, fernanda.watanabe, mauricio.galo}@unesp.br

Keywords: Mining detection, SAR, Sentinel-1, deep learning, environmental monitoring.

Abstract

Illegal mining represents a major environmental and socio-political threat in the Brazilian Amazon, particularly within protected areas and indigenous territories. While optical remote sensing has been widely used to detect mining activity, its utility is limited by persistent cloud cover. This paper explores the potential of C-band Synthetic Aperture Radar (SAR) imagery from Sentinel-1, combined with a lightweight convolutional neural network (CNN), to identify illegal mining under such challenging conditions. The model was trained on seven Sentinel-1 scenes from the Tapajós basin (state of Pará) and evaluated both within this training region and on an independent test set from the Yanomami Indigenous Territory (state of Roraima), using reference data from the Amazon Mining Watch (AMW) project. A total of 2,394 labelled patches supported supervised training. Results show balanced performance in the Tapajós region (F1-score = 0.676) and robust generalization to the Yanomami region (F1-score = 0.630. Most errors were associated with peripheral mining structures and small-scale disturbances, reflecting difficulties in capturing low-density mining patterns. Overall, the findings demonstrate the full potential of SAR-based deep learning approaches for monitoring illegal mining in persistently cloud-covered Amazonian landscapes. Future improvements may come from integrating terrain variables such as elevation and hydrological proximity, as mining often follows narrow streams (*igarapés*) closely tied to local topography.

1. INTRODUCTION

Illegal mining stands as one of the most urgent environmental challenges in Brazil, particularly within the Amazon region. Over the past years, an increase in mining requests has intensified pressure on protected ecosystems, with many operations starting under the pretext of small-scale extraction but rapidly expanding into broader, often unregulated activities (Ferreira et al., 2014; Villén-Pérez et al., 2018). This unchecked growth has increasingly encroached upon conservation units and indigenous territories, where oversight is often limited or absent.

The socio-environmental consequences of this expansion are severe. Invasions into territories belonging to indigenous groups such as the Munduruku and Yanomami have been closely linked to illegal miners and criminal networks seeking to exploit gold deposits in remote areas (Ferrante et al., 2021). These intrusions not only degrade biodiversity and contaminate water resources, but also incite violence and force the displacement of traditional communities. Unfortunately, safeguards agencies like FUNAI (National Indigenous Peoples Foundation) have suffered significant reductions in operational capacity, while the National Mining Agency continues to register increasing mining demands (Villén-Pérez, 2022). This institutional imbalance has left indigenous lands more vulnerable than ever to illegal occupation and exploitation.

Within this context, territories such as Kayapó and Tapajós have become emblematic of the broader crisis. Despite efforts by environmental enforcement agencies to dismantle illegal mining camps and seize machinery (MMA, 2023), or at least mitigating its effects, the vastness and inaccessibility of the terrain hinder consistent monitoring and control. Adding to the complexity, studies have revealed a strong association between illegal gold extraction and widespread environmental degradation. Much of this activity is believed to occur without formal registration, undermining official oversight and amplifying ecological risks (Manzolli et al., 2021).

In response, independent monitoring initiatives have filled critical gaps. Projects like MapBiomas have leveraged remote sensing and lightweight convolutional neural networks (CNNs) to map land use and detect illegal mining features using Sentinel-2 optical imagery (MapBiomas, 2023; Amazon Mining Watch, 2023). These models integrate and analyze multispectral data in small patches and aggregate imagery over multi-month periods to mitigate interference from cloud cover—a recurring limitation in the Amazon basin.

Even with these advancements, optical remote sensing remains constrained by the region's persistent cloudiness. In this regard, Synthetic Aperture Radar (SAR) has emerged as a powerful alternative. With its ability to penetrate clouds and operate under all weather and lighting conditions, SAR offers reliable and continuous image acquisition. Sentinel-1C, part of the Copernicus program, delivers high-resolution C-band RADAR imagery suitable for monitoring mining-related disturbances across the Amazon, regardless of seasonal visibility challenges (ESA, 2023; Jensen, 2011).

Considering this context, the aim of this paper is to adjust a convolutional neural network (CNN) model, originally designed for detecting large-scale mining areas using optical Sentinel-2 imagery, to detect smaller-scale illegal mining activities in the Amazon region. The research proposes leveraging Synthetic Aperture Radar (SAR) data from Sentinel-1 to overcome the limitations imposed by adverse weather conditions, such as frequent cloud cover and rainfall, which hinder optical imagery usage. By modifying the network architecture and integrating

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W4-2025-205-2026

SAR data, this work seeks to provide a more robust and efficient approach for monitoring illegal mining activities, contributing to environmental preservation and territorial security in hard-to-reach areas.

2. METHODOLOGY

To provide a general overview of the methodological steps, a flowchart was developed summarizing the main stages of the proposed approach (Figure 1). The workflow delineates the sequence from data acquisition and preprocessing to model training, classification, and validation, specifying the data acquisition platforms and computational tools utilized in processing. This schematic serves to clarify the integrated workflow employed for detecting illegal mining using Sentinel-1 SAR imagery.

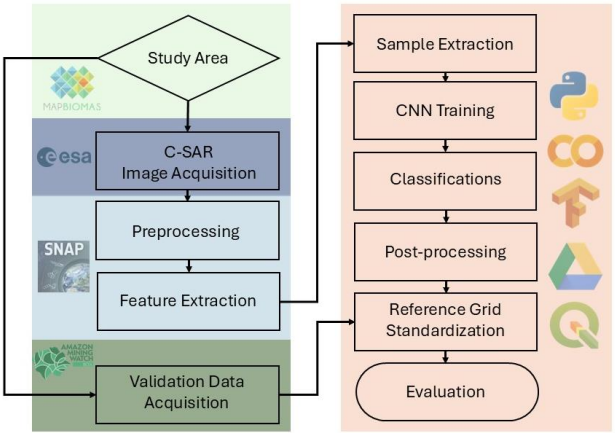

Figure 1. Overview of the methodological workflow for mining detection using Sentinel-1 SAR imagery.

2.1 Study areas

The study was conducted in two distinct regions of the Brazilian Amazon, both heavily affected by gold mining activities but with different socio-environmental contexts (Figure 2). The first area, used for model training, is located in the western portion of the state of Pará, within the Tapajós River basin. The selected region lies near, but outside, the legally demarcated Munduruku Indigenous Territory. This region has been historically impacted by widespread illegal mining activity, with observable land surface alterations and extensive use of machinery along riverbanks and forested areas. The availability of consistent SAR backscatter patterns, associated with mining activities, made this region suitable for training and validating the segmentation model.

In contrast, to evaluate the model's generalization capabilities, classifications were performed in a second area located on the opposite side of the Amazon, in the state of Roraima, along the northern frontier with Venezuela. This region encompasses part of the Yanomami Indigenous Territory, a remote and ecologically sensitive area that has faced a dramatic surge in illegal mining operations in recent years. Unlike the training region, mining activities here occur within indigenous land boundaries, presenting unique environmental, legal, and humanitarian concerns. By applying the trained model to this geographically and culturally distinct area, we aim to assess its potential for large-scale, cross-regional detection of illegal mining in varied Amazonian landscapes.

In this context, two classifications were generated. The first refers to applying the trained model to the Pará/Tapajós region, from which the training data was extracted and to which the model was specifically tailored. The second involves applying the model to the Roraima/Yanomami region, geographically and environmentally distinct from the training area, to evaluate its ability to generalize and detect illegal mining in previously unseen landscapes.

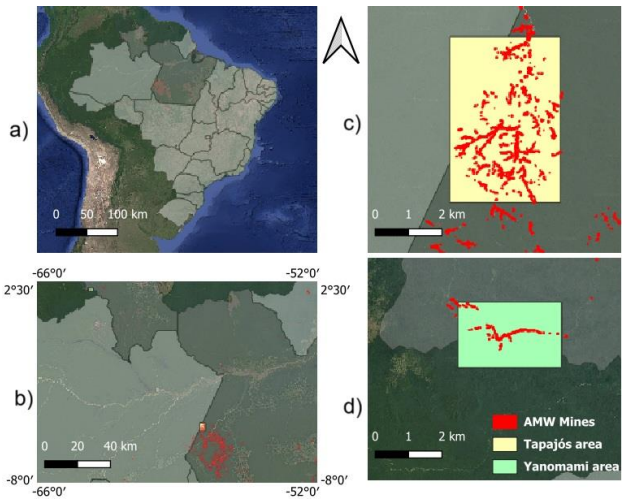

Figure 2. Localization of the areas, ground truth detections from Amazon Mining Watch (AMW) appear in red. (a, b) Study areas displayed on a Google Satellite background with an overlay of the Brazilian state boundaries. (c) Tapajós area, where training samples were extracted, is highlighted in yellow. (d) Yanomami area, used for generalization, is shown in green.

2.2 Image Acquisition

For this study, Sentinel-1 imagery was selected due to its all-weather imaging capabilities and independence from solar illumination, making it particularly suitable for applications in the Amazon region, where persistent cloud cover and rainfall frequently hinder optical remote sensing. All images were acquired in Interferometric Wide Swath (IW) mode, which is the default acquisition mode for land applications. IW mode collects dual-polarized (VV and VH) Ground Range Detected (GRD) products over a 250 km swath using the TOPSAR (Terrain Observation with Progressive Scans SAR) technique, with an approximate spatial resolution of 10 m.

Data were accessed through the European Space Agency's Copernicus Open Access Hub (https://scihub.copernicus.eu/) under the Copernicus free and open data policy. In the Tapajós basin (Pará), multiple Sentinel-1 scenes acquired throughout 2023 (January 23, June 16, August 3, September 8, October 2, November 7, and December 25) were used to extract training samples. Classification in this region was performed using a scene acquired on December 25, 2023, ensuring consistency with the samples extracted throughout the year. These dates were chosen to capture different seasonal and activity conditions in areas affected by mining.

The Yanomami Indigenous Territory was classified using a Sentinel-1 scene acquired on January 2, 2024. Since the validation grid is based on the 2023 AMW ensemble, gathering all detections that occur throughout the year. The short temporal gap ensures comparability, as mining areas are unlikely to have changed substantially between late 2023 and early 2024. A

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W4-2025-205-2026 | © Author(s) 2026. CC BY 4.0 License.

targeted spatial subset of the Yanomami region was defined for classification, focusing on zones with prior mining reports or visual indicators, in order to evaluate the model's generalization beyond the training domain.

2.3 Preprocessing

All preprocessing procedures were performed using ESA's Sentinel Application Platform (SNAP), following the standard workflow for Ground Range Detected (GRD) products. The steps included the application of precise orbit files to enhance geolocation accuracy, removal of thermal noise, and radiometric calibration to convert digital numbers into backscatter coefficients (σ^0) expressed in decibels. Speckle noise was attenuated using the Lee filter with a 3×3 window, and terrain correction was carried out using the SRTM Digital Elevation Model (DEM) with 1" (~ 30 m) to project the imagery from slant to ground range, ensuring accurate geocoding.

As part of this process, spatial subsets were manually delineated using shapefiles for both the training and classification regions. These subsets were designed to focus on areas with observed or reported illegal mining activity, allowing the study to concentrate on relevant targets while reducing computational complexity. This localized approach reflects the intent to assess the applicability of SAR-based deep learning methods, within discrete scenario-specific contexts, rather than aiming for generalized regional classification.

Following preprocessing, each Sentinel-1 scene was structured into three layers: σ^0_{VV}, σ^0_{VH}, and the Polarization Averaging (PA) index, calculated from the calibrated dual-polarized data ($(\sigma^0_{VH} - \sigma^0_{VV})/2$). Among several tested polarimetric indices, the PA index was selected due to its superior visual sensitivity to spatial patterns commonly associated with mining activity. These attributes formed the basis for the extraction of square image patches centered on reference points, which were used to build the dataset for training (Figure 3) and validating the convolutional neural network (CNN) model.

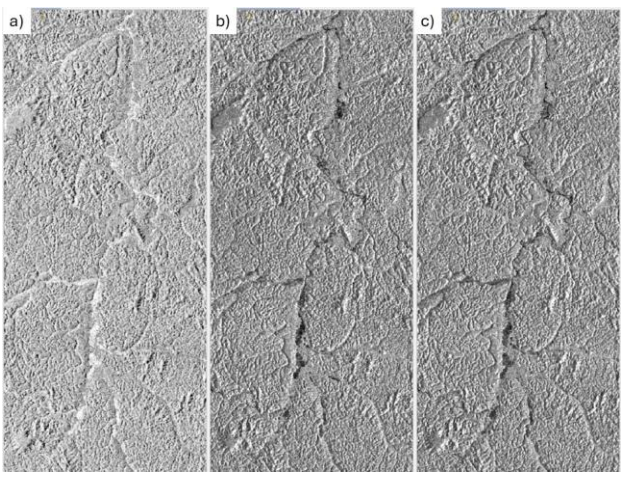

Figure 3. C-SAR polarimetric features images of the training area (Tapajós), acquired on January 23, 2013: (a) PA index, (b) σ^0_{VH}, (c) σ^0_{VV}.

2.4 Sample extraction

The preparation of the training dataset consisted of selecting square image patches centered on manually defined reference points within the Tapajós region. Initially, the patch size of 520 × 520 m was considered in the validation step, but it produced suboptimal results for this application. The patch dimension was therefore increased to 1,040 × 1,040 m, resulting in four times the area per sample. This adjustment enhanced the access to contextual information available to the model, which is particularly important in mining detection scenarios where spatial patterns may extend beyond smaller windows.

Detections from the Amazon Mining Watch (AMW) project were employed as external reference data to support the acquisition of training samples. In the Tapajós area, AMW detections guided the extraction of mining and non-mining samples, ensuring precise localization of patches for both classes. In contrast, in the Yanomami region, AMW detections were used exclusively as independent reference data for model evaluation and validation.

In the Tapajós training region, 171 reference points were delineated for each class (mining and non-mining) using QGIS. These annotations were based on visual interpretation of Sentinel-2 imagery (2023) and Sentinel-1 backscatter signals, and were cross-validated against detections from Amazon Mining Watch (AMW). Mining areas were typically identified as expansions along water bodies, marked by irregular textures, persistent clearings, and strong backscatter contrasts, pattern commonly linked to deforestation and machinery activity. In contrast, non-mining areas displayed vegetation-like backscatter patterns with smoother spatial transitions. Figure 4 presents representative annotated patches for both classes

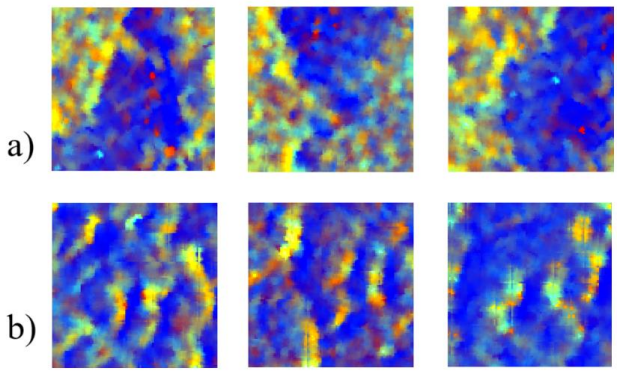

Figure 4. Representative patches utilized in model training, derived from the Sentinel-1 scene captured on August 3, 2023. The RGB composition is defined as: R = σ^0_{VV}, G = σ^0_{VH}, B = PA. (a) Mining areas are distinguished by continuous blue regions. (b) Non-mining areas exhibit more heterogeneous patterns, where the bluish regions are fragmentated.

Each point was used to extract a single image patch from each of the seven Sentinel-1 scenes used for training, resulting in a total of 1,197 patches (see examples in Figure 4) per class (171 points × 7 scenes × 2 output classes). The acquisitions were automated via a Python script developed in *Google Colab*, which reads multipoint shapefiles and executes spatial windowing operations with *Rasterio* command. The script also handles validation of sample dimensions, manages edge conditions, and stores the final patches as GeoTIFF files, with updated georeferencing metadata. This approach ensured spatial consistency across scenes and preserved class balance in the dataset, providing a robust foundation for supervised training of the CNN model.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-X-3-W4-2025-205-2026

2.5 Convolutional Neural Network Architecture and Training

The convolutional neural network (CNN) model was implemented using *TensorFlow/Keras* for binary classification of mining and non-mining areas based on input patches with three channels (σ^0_{VV}, σ^0_{VH}, and PA index). The input tensor size is 104 × 104 × 3.

The CNN architecture consists of three convolutional blocks followed by a global average pooling layer and dense layers. Specifically:

- First convolutional block: Conv2D with 32 filters, 3×3 kernel, ReLU activation, and same padding → BatchNormalization → MaxPooling2D with 2×2 pool size.
- Second convolutional block: Conv2D with 64 filters, 3×3 kernel, ReLU activation, and same padding → BatchNormalization → MaxPooling2D with 2×2 pool size.
- Third convolutional block: Conv2D with 128 filters, 3×3 kernel, ReLU activation, and same padding → BatchNormalization → GlobalAveragePooling2D.

A Dropout layer with a probability of 0.5 is applied after the global pooling, followed by a dense layer with 64 neurons and ReLU activation. The final output layer has 2 neurons with softmax activation for binary classification. The key hyperparameters and network configurations are summarized in Table 1.

Configuration / Hyperparameter	Specification / Value
Input size	104 x 104 x 3
Conv layers	3 layers: 32→64→128 filters, 3×3 kernel, ReLU, same padding
Pooling layers	MaxPooling 2×2 after first 2 conv layers
Global pooling	GlobalAveragePooling2D
Batch normalization	After each conv layer
Dropout	0.5 (after global pooling)
Dense layers	64 neurons, ReLU
Output layer	2 neurons, softmax
Optimizer	Adam
Learning rate	0.001
β1	0.9
β2	0.999
Epsilon	1e-8
Loss function	Focal loss (γ=2.0, α=0.25)
Batch size	32
Epochs	30
Validation split	0.2

Table 1. Hyperparameters used in the CNN training.

To address class imbalance, the focal loss function (Lin et al., 2020) was employed with $\gamma = 2.0$ and $\alpha = 0.25$. Focal loss dynamically scales the standard cross-entropy loss by reducing the relative loss contribution of well-classified examples, thereby focusing the training process on harder, misclassified samples. This is particularly important in illegal mining detection, where non-mining areas vastly outnumber mining patches, ensuring that the model does not become biased toward the majority class.

The network was trained using the Adam optimizer with learning rate 0.001, $\beta1 = 0.9$, $\beta2 = 0.999$, and $\varepsilon = 1e\text{-}8$. Training was conducted for 30 epochs with a batch size of 32, and the dataset was split into training and validation subsets (80/20). Input data were standardized using z-score normalization.

2.6 Classifications and Post-processing

The trained model was applied to two distinct regions using pre-processed Sentinel-1 imagery: (i) the training region in the Tapajós basin (Pará), and (ii) a generalization region in the Yanomami Indigenous Territory (Roraima). The Tapajós classification used a scene acquired on December 25, 2023, while the Yanomami scene was from January 2, 2024. This dual-region setup enabled the evaluation of both within-domain performance and out-of-sample generalization.

Classification was performed using a sliding window approach with a stride of 52 pixels (50% patch overlap) (Figure 5a). Patches classified as "mines" with confidence above the threshold were retained, and the centroid of each positive patch was converted to geographic coordinates (latitude and longitude).

A spatial adjustment step was applied to generate polygon outputs compatible with the validation dataset and suitable for geospatial visualization. Each centroid was expanded into four adjacent square polygons (520 × 520 m), creating a square around the detected point (Figure 5b). This procedure was specifically designed to align the size of the classification outputs with the AMW ground truth polygons, ensuring spatial comparability during validation. Care was taken to preserve the positional accuracy of the points during the coordinate transformations.

A spatial thresholding procedure was applied to remove redundant polygons generated by the sliding-window approach. Since overlapping patches can produce polygons derived from adjacent centroids with substantial spatial overlap, polygons sharing more than 90% of their area were merged by retaining a single representative instance. Although the stride was set to 50%, minor discrepancies may still arise; thus, any pair of squares with at least 90% overlap is merged into a single one. (Figure 5c). Figure 5 develops an example, on Figure 5b there are eight squares (520 x 520 m), where four squares are overlapped, after adjustment, the overlapped squares become two squares without overlap.

The final outputs were saved as GeoJSON files, producing spatially consistent and non-redundant polygon layers ready for visual inspection and comparative analysis.

The classification process was repeated twice, using CNN classification confidence thresholds of 0.80 and 0.90, resulting in four mining detection maps ~~classification runs~~. Higher thresholds were selected instead of the conventional 0.5 to reduce false positives and increase the reliability of detected mining areas. This choice prioritizes precision over recall, which is particularly important in the context of mining detection, where misclassified areas could lead to incorrect environmental assessments or enforcement decisions.

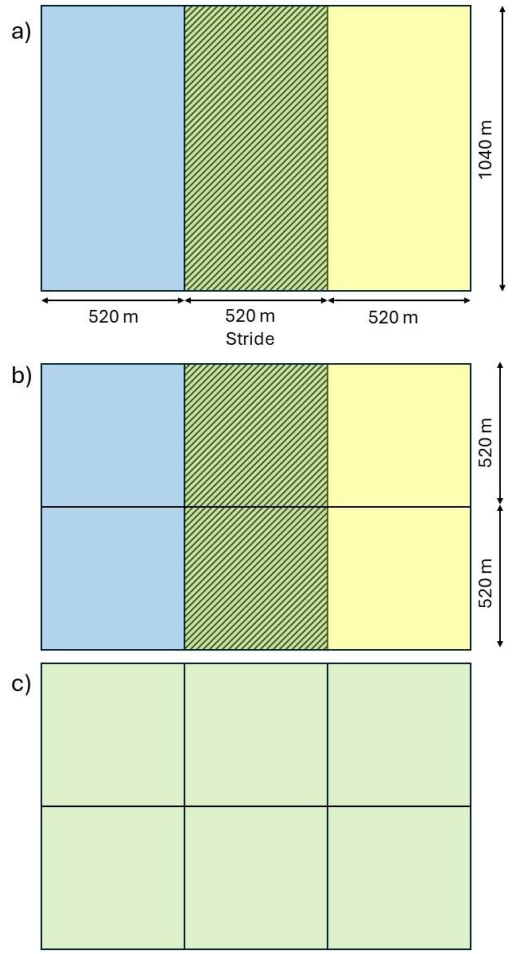

Figure 5. Spatial adjustment procedure to eliminate overlapping area between patches. (5a) Two 1040 × 1040 m squares with 50% overlap. (5b) Each square is subdivided into four 520 × 520 m cells. (5c) Overlapping squares merged into a single representative polygon.

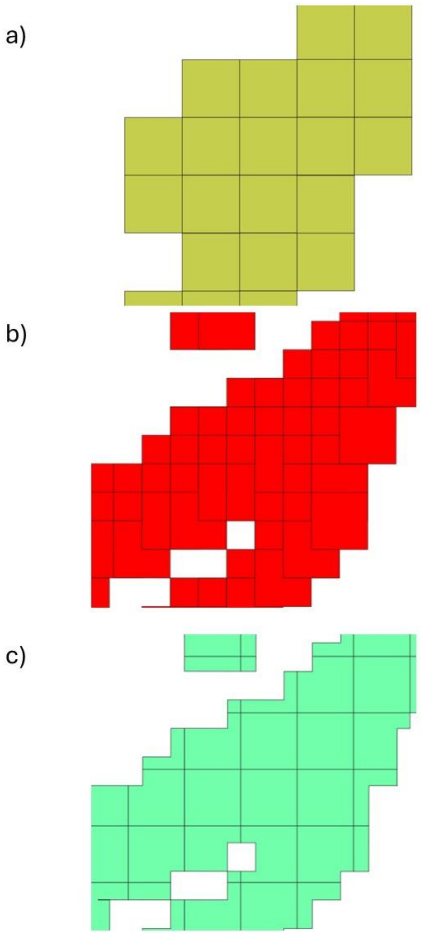

Figure 6. Compatibilization of classification and validation datasets. (a) Model classification output represented as a regular grid. (b) Validation dataset derived from AMW detections, containing overlapping polygons. (c) Standardized validation grid, standardized into a non-overlapping structure for consistent comparison with model outputs.

2.7 Reference Grid Standardization

The reference dataset (AMW detections) used for validation was derived from an ensemble of mining classifications across multiple Sentinel-2 scenes acquired in 2023. This cumulative process produced overlapping polygons, leading to inconsistencies in the spatial structure of the validation mask. Unlike training, which uses labelled sample patches, validation requires a standardized spatial grid to enable consistent and meaningful comparison with model classifications.

To adjust the dataset, a spatial standardization step was applied. All polygons were dissolved into a single geometry and reprojected to a metric coordinate system. A uniform grid of square cells measuring 520 × 520 m was generated over the area of interest. The grid was clipped to the bounds of the reference geometry, and small fragments covering less than 20% of a full cell (<10,816 m²) were discarded, ensuring that only meaningful portions of the reference area were retained.

This process produced a non-overlapping, uniformly structured validation grid compatible with the CNN model outputs, enabling reliable spatial comparisons and post-classification evaluation (Figure 6).

2.8 Classification performance evaluation

To evaluate the performance of the CNN model in detecting illegal mining areas, standard binary classification metrics based on spatial intersections between classification results and ground truth were computed. Each patch was assigned to one of the following categories:

-True Positive (TP): a patch classified as mining that overlaps a reference mining polygon.
-False Positive (FP): a patch classified as mining that does not overlap any reference mining polygon.
-False Negative (FN): a reference mining polygon that does not overlap any patch classified as mining.
-True Negative (TN): patches correctly classified as non-mining were not explicitly computed, as the vast majority of the Amazon region corresponds to non-mining areas, which would dominate the metrics and reduce the sensitivity to mining detection.

Based on these counts, the following evaluation metrics were calculated:

- Precision - quantifies the proportion of correctly identified mining patches among all patches classified as mining (eq. 1).
- Recall - measures the proportion of actual mining areas that were correctly detected by the model (eq. 2).

- F1-score - represents the harmonic mean of precision and recall, providing a single measure that balances both types of errors (eq. 3).

$$\text{Precision} = (TP) / (TP+FP) \quad (1)$$

$$\text{Recall} = (TP) / (TP+FN) \quad (2)$$

$$\text{F1-score} = (2 * \text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall}) \quad (3)$$

The CNN model's performance was assessed by comparing its classifications to a validation grid based on ensemble mining data (AMW datasets) from multiple Sentinel-1 scenes acquired in 2023. The evaluation was based on spatial intersection between classified and reference polygons, using a logical rule set to label each patch as true positive (TP), false positive (FP), or false negative (FN). Patches were labelled as TP if their polygons intersected any polygon in the validation grid. Polygons with no intersections were labelled as FP, while reference polygons with no overlapping classifications were labelled as FN. True negatives (TN) were not explicitly computed, as the vast majority of the Amazon region corresponds to non-mining areas. Including TN values would dominate the evaluation metrics and reduce sensitivity to model performance in detecting actual mining areas.

This evaluation approach emphasizes spatial coverage over pixel-level accuracy, identifying potentially affected zones rather than exact boundaries. It was applied to both the Tapajós region (training domain) and the Yanomami region (generalization domain), at confidence thresholds of 0.80 and 0.90. Each patch was labelled according to its classification and exported as a GeoJSON file for spatial visualization.

Figures 7 and 8 illustrate the model classifications, where panels a1 and a2 correspond to mining areas identified in the Tapajós and Yanomami regions, respectively.

3. RESULTS AND DISCUSSIONS

Figure 7 presents the classification results, in terms of detected mining areas, for the Pará/Tapajós training region using two different confidence thresholds: 0.80 and 0.90. Additionally, model performance is illustrated through the spatial representation of polygons categorized as true positives (TP), false positives (FP), and false negatives (FN). Figure 8 presents the corresponding classifications for the Roraima/Yanomami test area, applying the same thresholds to assess the model's generalization capability and highlight its successes and misclassifications.

Figures 7 (Tapajós area) and 8 (Yanomami region) present evaluation metrics in a spatial format that offers an intuitive framework for assessing the model's mining detection performance. This representation facilitates comparative analysis across different confidence thresholds and geographic regions.

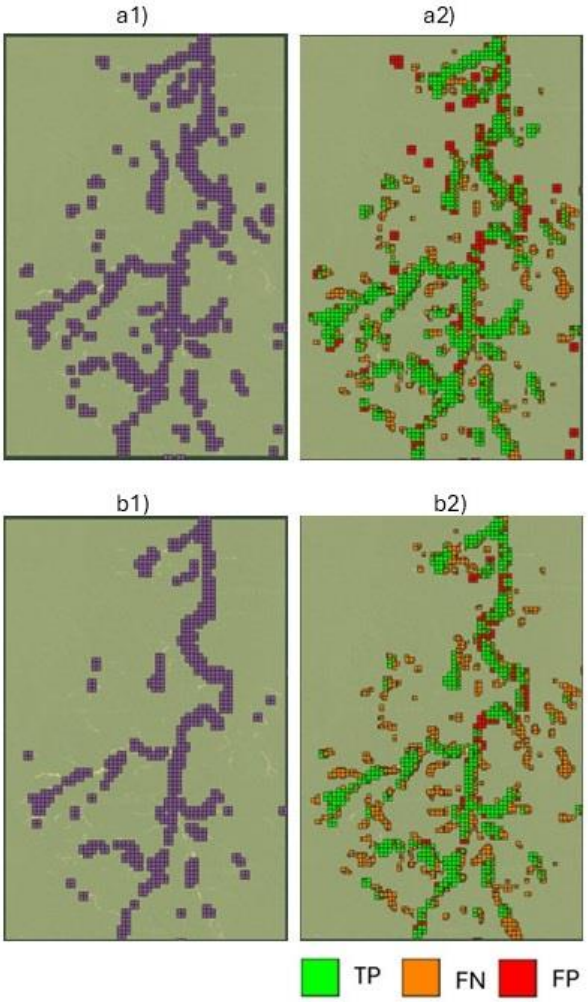

Figure 7. Classification results and model evaluation for the Tapajós region. Mining areas detected at thresholds of 0.8 (a1) and 0.9 (b1) are shown in purple color. Spatial distribution of detections coinciding with the AMW reference and false detections are expressed in terms of TP, FN and FP, at thresholds of 0.8 (a2) and 0.9 (b2).

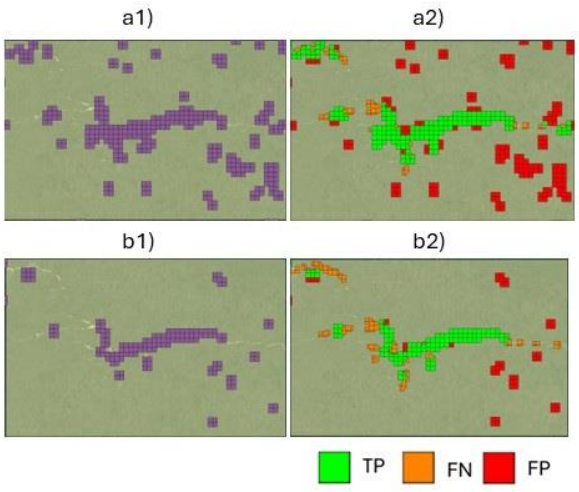

Figure 8. Classification results and model evaluation for the Yanomami test area. Mining areas detected at thresholds of 0.8 (a1) and 0.9 (b1) are shown in purple colour. Spatial distribution

of detections coinciding with the AMW reference and false detections are expressed in terms of TP, FN and FP, at thresholds of 0.8 (a2) and 0.9 (b2).

In both regions, a considerable number of FN were observed, particularly associated with mining extensions and peripheral branches of known exploitation zones. These features tend to be more fragmented and cover a smaller fraction of each patch, reducing their prominence in the RADAR backscatter signal. Given the spatial resolution of the Sentinel-1 data and the patch-based classification strategy, mining activity that is dispersed, newly established, or embedded within dense forest cover may not produce a strong enough signal to surpass the classification threshold. This highlights a limitation of the current model in detecting low-density or marginal mining activity, and suggests that incorporating higher-resolution data or temporal aggregation strategies may help mitigate this issue in future implementations.

Table 2 presents the quantitative results for all four classification configurations, highlighting key detection quality metrics (Recall, Precision, and F1-Score) calculated from TP, FP, and FN values. Recall measures the model's sensitivity, Precision reflects confidence in detections, and F1-Score represents the harmonic mean of Recall and Precision.

Region	Threshold	Precision	Recall	F1-Score
Tapajós	0.80	0.805	0.583	0.676
Tapajós	0.90	0.882	0.430	0.578
Yanomami	0.80	0.476	0.775	0.589
Yanomami	0.90	0.710	0.567	0.630

Table 2. Evaluation metrics for classifications at different confidence thresholds.

As expected, increasing the confidence threshold from 0.80 to 0.90 led to a significant improvement in precision, at the expense of lower recall. This trade-off was consistent across both regions. In the Tapajós area (training domain), the model achieved higher overall performance due to exposure to similar spatial and contextual patterns during training. Conversely, the Yanomami region, which was not included in the training set, presented a different morphological and spectral configuration. This likely contributed to the higher FP rate observed at the lower threshold. The improved F1-score at 0.90 in the Yanomami regions suggests that a more conservative classification strategy may be required in unfamiliar regions, where the model's generalization ability is challenged.

4. CONCLUSIONS

This paper explored the potential of using Sentinel-1 SAR imagery combined with a convolutional neural network (CNN) for the detection of illegal mining activity in the Brazilian Amazon. The approach demonstrated promising results in the Tapajós region, where the training data were sourced, and showed moderate generalization capacity when applied to the Yanomami Indigenous Territory, an unseen and morphologically distinct region.

The results highlight the importance of confidence threshold selection: while higher thresholds reduced false positives, they also decreased recall, particularly in areas where mining activity is more fragmented or spatially diluted. A significant number of FN were associated with peripheral or branching mining structures, which likely reflect the RADAR sensor's limited sensitivity to small-scale or low-density disturbances.

The methodology proved effective for mapping high-confidence mining zones and offers a viable alternative to optical-based monitoring, especially in cloud-covered regions. However, further developments are needed to improve detection in marginal cases. Overall, the findings support the feasibility of SAR-based deep learning methods as a complementary tool for environmental monitoring in the Amazon region, with potential applications in enforcement, early warning, and policy support.

Future work may include expanding the training dataset to cover more diverse mining scenarios, incorporating temporal data fusion, and testing architectures with finer spatial granularity or attention mechanisms to enhance sensitivity to dispersed activity patterns. Since topographic variables such as elevation, slope, and hydrological context can offer valuable information, as illegal mining activities frequently occur along small watercourses, such as igarapés, the integration of DTM (Digital Terrain Models) and DSM (Digital Surface Model) can also contribute. Another alternative for further investigation is the use of deep learning architectures based on Transformers, which have shown promising results in remote sensing tasks by capturing long-range spatial dependencies and contextual relationships more effectively than conventional CNNs.

Acknowledgments

The authors are grateful for the support provided by the Graduate Program in Cartographic Sciences (PPGCC) and the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brazil (CAPES).

References

Amazon Mining Watch, 2023: Metodologia. https://amazonminingwatch.org/pt/methodology (29 April 2025).

Balaniuk, R., Isupova, O., Reece, S., 2020: Mining and tailings dam detection in satellite imagery using deep learning. *Sensors*, 20(23), 6936. https://doi.org/10.3390/s20236936.

Ferrante, L., Fearnside, P.M., 2021: Brazilian government violates Indigenous rights: What could induce a change? *DIE ERDE – J. Geograph. Soc. Berlin*, 152(3), 200–211. https://doi.org/10.12854/erde-2021-584.

Ferreira, J., et al., 2014: Brazil's environmental leadership at risk. *Science*, 346(6210), 706–707. https://doi.org/10.1126/science.1260194.

Jensen, J.R., 2011: *Sensoriamento remoto do ambiente: uma perspectiva em recursos terrestres*. 2nd ed., Parêntese, São José dos Campos.

T. -Y. Lin, P. Goyal, R. Girshick, K. He and P. Dollár, 2020: Focal Loss for Dense Object Detection. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 42, no. 2, pp. 318-327, 1 Feb. 2020, doi: 10.1109/TPAMI.2018.2858826.

Ma, L., Liu, Y., Zhang, X., et al., 2019: Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS J. Photogramm. Remote Sens.*, 152, 166–177. https://doi.org/10.1016/j.isprsjprs.2019.04.015.

MapBiomas, 2023: Coleção 5.0 dos mapas anuais de cobertura e uso da terra do Brasil. https://brasil.mapbiomas.org/ (11 April 2025).

Manzolli, B.A., Rajão, R., Bragança, A.C.H., Oliveira, P.T.M., 2021: Legalidade da produção de ouro no Brasil. UFMG/MPF, Belo Horizonte.

Ministério da Ciência, Tecnologia e Inovações (MCTI), 2023: Garimpo ilegal e o avanço da prática na Terra Indígena Kayapó. https://www.gov.br/museugoeldi/pt-br/arquivos/noticias/garimpo-ilegal-e-o-avanco-da-pratica-na-terra-indigena-kayapo (29 April 2025).

Ministério do Meio Ambiente (MMA), 2023: Equipe do IBAMA é atacada a tiros em ação de combate ao garimpo ilegal na Bacia do Tapajós. https://www.gov.br/mma/pt-br/equipe-do-ibama-e-atacada-a-tiros-em-acao-de-combate-ao-garimpo-ilegal-na-bacia-do-tapajos (29 April 2025).

Ministério das Minas e Energia, 1973–1985: Projeto RADAMBRASIL. *Levantamento de recursos naturais*, Vols. 1–44. Ministério das Minas e Energia, Rio de Janeiro.

Veríssimo, A., et al., 2011: Protected areas in the Brazilian Amazon: challenges and opportunities. Imazon, Belém and Instituto Socioambiental (ISA), São Paulo.

Villén-Pérez, S., et al., 2018: Mining code changes undermine biodiversity conservation in Brazil. *Environ. Conserv.*, 45(1), 96–99. https://doi.org/10.1017/S0376892917000376.

Villén-Pérez, S., et al., 2022: Mining threatens isolated indigenous peoples in the Brazilian Amazon. *Global Environ. Change*, 72, 102398. https://doi.org/10.1016/j.gloenvcha.2021.102398.