Evaluation of a YOLOv8-Based Method for Detecting Unauthorized Airstrips in the Amazon Rainforest using SAR Imagery

Leandro da Silva Gomes, Elcio Hideiti Shiguemori, Tahisa Neitzel Kuck, and Dimas Irion Alves

***Abstract*— This study presents the application of a method to detect unauthorized airstrips in the Brazilian Amazon using synthetic aperture radar (SAR) images from the Sentinel-1 satellite. There are illegal activities, such as mining and drug trafficking, which use these airstrips for their operations. Furthermore, the area is known for its dense vegetation. A YOLOv8 model was trained on a database comprising 646 training, 277 validation, and 117 generalization testing images. Results demonstrate the model's effectiveness in identifying unauthorized airstrips, highlighting its potential for monitoring remote regions. The central experiment correctly identified 46 out of 117 landing strips.**

***Keywords*— SAR images, illegal airstrips, Amazon Rainforest, Sentinel-1, YOLO**

I. Introduction

The Brazilian Amazon, covering over 5 million square kilometers, is a vast and diverse territory rich in biodiversity and natural resources [1]. However, this region faces a series of challenges related to illicit practices. A growing concern is the presence of unauthorized airstrips, which are frequently used for illegal activities such as mining and drug trafficking [2], [3]. Moreover, they are difficult to detect and monitor using visible spectrum satellite images due to the dense and constant cloud cover in the Amazon during certain periods of the year. This high prevalence of cloud cover complicates detecting illicit activities and assessing of environmental changes [4]. In this context, Synthetic Aperture Radars (SAR) has been a tool capable of overcoming such challenges by providing high-resolution images of terrestrial and maritime surfaces, irrespective of weather conditions and sunlight. The effectiveness of SAR radars in detecting targets and environmental changes, even in extensive areas and under various atmospheric conditions, is documented and recognized [4]–[9]. Figure 1 presents a SAR image obtained even under adverse weather conditions.

However, the complexity and multifaceted nature of these images can pose a significant challenge for human interpretation [6]. Given this complexity, integrating computational resources, such as Artificial Intelligence (AI), has proven to be a solution for extracting information from SAR images. With its capabilities for deep learning and processing large volumes of data, AI provides support to expedite decision-making using these images [4]–[7], [10], [11]. Some results presented in studies involving the automatic detection of targets in radar images point to a promising scenario for the use of artificial intelligence for this purpose [5], [6], [10], [11].

The thesis proposed in [6], for instance, introduces the application of a Convolutional Neural Network (CNN) for the automatic detection of small vessels in SAR images acquired by the ICEYE constellation. The central motivation of the work is to enhance vessel detection methods in scenarios where traditional visual techniques are limited, using the YOLOv4 network, adapted for identifying vessels in SAR images with a spatial resolution of 2.5 meters per pixel. The central experiment achieved a Mean Average Precision (mAP) of 64% and a global accuracy of 82.75% in generalization testing. However, it was observed that the developed algorithm exhibited a high rate of false alarms in high sea regions not included in the training, indicating the need for more data to refine the model and its generalization capability.

Suresh et al. [12] propose one of the first methods for automatically detecting oil spills in SAR images from the ENVISAT satellite. In this study, one observation was the dark tone displayed by these patches in the images due to diffraction effects occurring in backscatter. The automatic detection methods were able to distinguish oil spills from other similarly dark structures, achieving up to 76% accuracy in identifying the targets of interest.

Regarding the using a lighter YOLO model for object detection, Pang et al. [5] propose the YOLOv5-MNE for ship detection in SAR images, aiming to overcome the challenges posed by unclear outlines of targets and complex image backgrounds. The proposed solution is to be efficient in terms of training speed, execution memory, and model parameter count, maintaining acceptable accuracy in large datasets. The study introduces the MNEBlock module and replaces the Sigmoid Linear Unit (SiLU) activation function with the Rectified Linear Unit (ReLU) at the network's output. The intent was to reduce the number of parameters and incorporate the Coordinated Attention (CA) mechanism to enhance detection performance. The results show that the YOLOv5-MNE algorithm achieves a precision of up to 94.8%. The research also addresses studies to assess the impact of different modules and configurations on model performance, indicating the effectiveness of the MNEBlock and the CA mechanism in improving the precision of ship detection in SAR images.

Leandro Gomes and Dimas Irion are with the Aeronautics Institute of Technology (ITA), São José dos Campos-SP, Brazil, e-mail: leandro-gomes.vega@gmail.com, dimasirion@ita.br. Elcio Shiguemori and Tahisa Kuck are with the Institute for Advanced Studies (IEAv), São José dos Campos-SP, Brazil, email: elcio@ieav.cta.br, tatakuck@gmail.com.

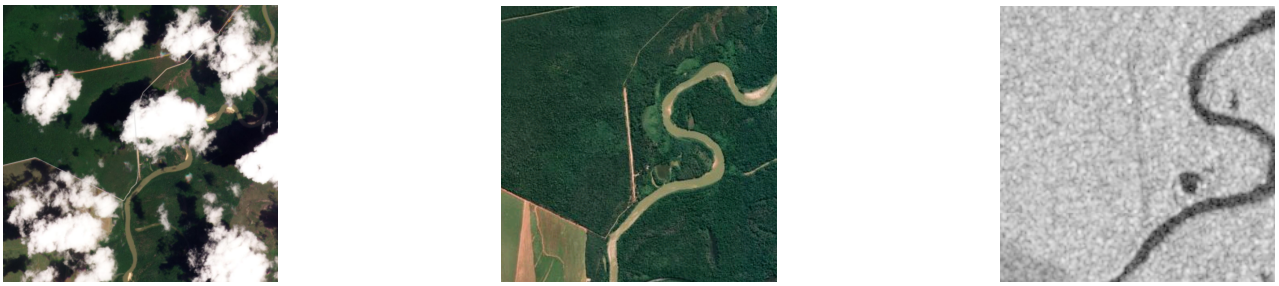

(a) PLANET image with clouds

(b) PLANET image without clouds

(c) Sentinel-1 SAR image VV polarization

Fig. 1: Airstrip near Culuene river, Brazil. Lon: -52.8774 Lat: -12.9757

However, no studies proposing the detection of unauthorized airstrips in the Amazon using SAR images were found.

Given the negative impacts caused by clandestine tracks in the Amazon, the advantages of using SAR images, the promising results from AI studies with the YOLO model, and the lack of research on detecting these tracks using SAR, this study aims to evaluate the performance of the YOLOv8 model in detecting such targets. [1]

This paper is structured as follows: Section 2 discusses the technical concepts, characteristics, and functionalities of the YOLOv8 architecture. Section 3, Materials and Methods, outlines the methodological procedures, including data collection, neural network configuration, training, and the metrics used for evaluation. Section 4 presents the study's findings, analyzing the performance of the YOLOv8 network in detecting unauthorized airstrips. Finally, the main conclusions and suggestions for future research are provided in the Final Remarks.

II. YOLOv8

Object detection is a key area of study in computer vision, requiring accurate and efficient models, with the YOLO (You Only Look Once) version 8 emerging as one of the most advanced for this task. This version was created by the research team at Ultralytics and brings significant improvements in object detection in images through the integration of various components and techniques. For example, the convolutional block that combines a 2D convolutional layer, 2D batch normalization, and a SiLU activation function, optimizing model execution. Next, the "C to F" block in the architecture divides the resulting feature maps for further processing, while the bottleneck block incorporates shortcuts for training efficiency and accuracy. The Spatial Pyramid Pooling Fast (SPPF) block, a modification of spatial pyramid pooling for increased speed, adds convolutional blocks and 2D max pooling layers, emphasizing the importance of treating objects of various sizes without loss of spatial information [13]–[15]. Figure 2 was generated from reference [16], which presents a summary of the structure of the YOLOv8 network.

Within the YOLOv8 architecture, Ultralytics developed five model variants, to meet different performance and computational resource requirements. These are YOLOv8n (nano),

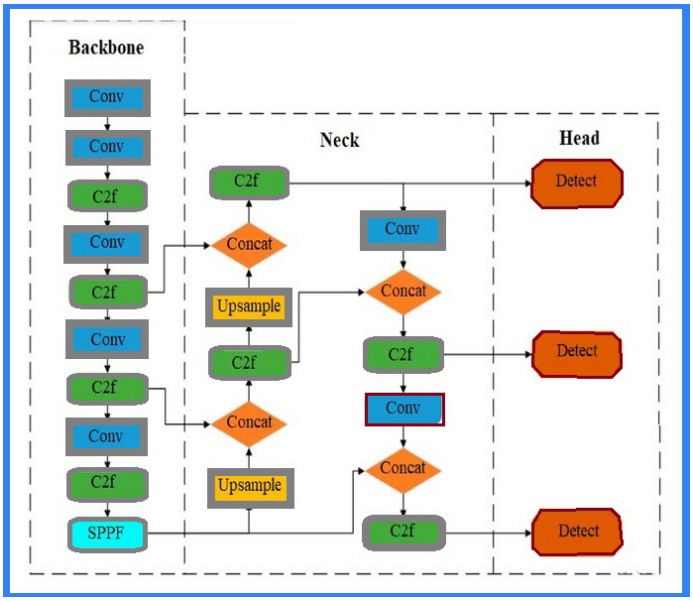

Fig. 2: Architecture of YOLO version 8.

YOLOv8s (small), YOLOv8m (medium), YOLOv8l (large), and YOLOv8x (extra large). Each variant is designed to balance precision and speed to be applicable across a wide range of devices, from those with limited computational capabilities, such as mobile devices, to highly robust systems. The YOLOv8n is the lightest and fastest, ideal for real-time applications with severe hardware constraints. Progressively, YOLOv8s, YOLOv8m, YOLOv8l, and YOLOv8x offer an increase in precision with a proportional cost in computational performance [17]. This version adopts an anchor-free model, making predictions directly on grid cells. The detection block, containing two paths for bounding box and class prediction, optimizes the accuracy of predictions using a sequence of convolutional blocks and a single 2D convolutional layer [17].

Overall, the YOLOv8 architecture is distinguished by its division into three main parts: Backbone, Bottleneck, and Head. The Backbone functions as a feature extractor, the Bottleneck effectively combines these features, and the structure's part called the Head makes the final predictions of classes and bounding boxes. The simplification in the integration between the Backbone and Head reflects a more direct and efficient

[1]Part of the results presented in this article was submitted to compete for the Mapbiomas Award. This award does not imply publication. Therefore, we consider that it does not constitute a double submission.

approach to object detection, consolidating YOLOv8 as a tool to tackle complex challenges in computer vision [17].

III. Materials and Methods

Initially, the position of the airstrips was verified in MapBiomas, a multi-stakeholder collaboration involving academia, technology firms, and NGOs in Brazil, focused on mapping and monitoring changes in land use and cover nationwide using advanced remote sensing and geospatial analysis. The project provides geographic data, including locations of authorized and unauthorized landing strips, assisting in logistical planning and field studies in the Amazon. Figure 3 shows the location of the airstrips mapped by the project.

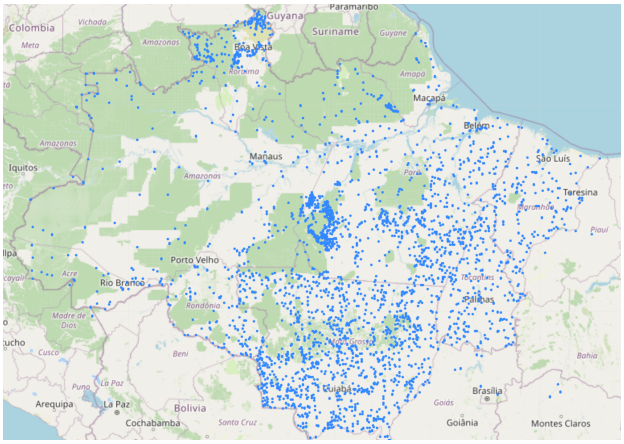

Fig. 3: Airstrips in the Amazon Rainforest recorded by MapBiomas.

The list of geographical coordinates, expressed in decimal degrees of longitude and latitude, was compiled to identify locations with unauthorized airstrips. This list was incorporated into a Python script executed in the Google Colab environment. The initial script procedure involved installing and configuring the Google Earth Engine (GEE) API, providing full access to the platform's data manipulation and geospatial analysis capabilities. For each location identified as containing an airstrip, an image slice was taken, covering an area of 4 km square in 200 × 200-pixel images from the Sentinel-1 satellite, corresponding to the year 2023 (the year of the last update of the MapBiomas airstrip database). The images were obtained in the Interferometric Wide (IW) mode and vertical-vertical (VV) polarization. Figure 4 exemplifies two images from the dataset.

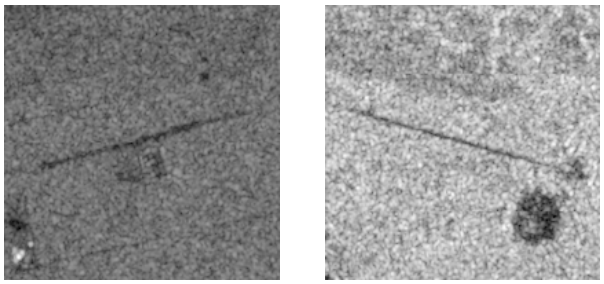

Fig. 4: Examples of images generated for the dataset.

Subsequently, the selected images were exported for storage in the cloud, followed by downloading for processing in a local environment. To prepare the images for analysis using the YOLOv8 neural network, the original data was converted to the PNG format using 8-bit integers. The images were labeled in the standard required by YOLOv8, using bounding boxes created directly on the SAR images with Python, aided by the use of high-resolution images from PLANET, corresponding to the exact dates as the SAR images. This approach was adopted due to the need for accurate identification of the airstrips, as the visual characteristics of these structures often appear indiscernible in radar images. Moreover, the targets are volatile: the airstrips can be temporary, with vegetation quickly covering previously deforested areas, or they can emerge suddenly due to illicit activities such as mining, requiring checks in visible spectrum images for their existence to ensure correct labeling. At the end of the image acquisition process, 1040 airstrips were obtained. The targets were randomly divided as follows: 646 images for training, 277 for validation, and 117 for generalization testing. The pre-trained YOLOv8x model, chosen for this study, has the most robust architecture in the series, with 68.2 million parameters, making it particularly suitable for scenarios that require high precision and superior performance. This model was trained using a 16GB V100 GPU to handle large datasets and high computational complexity.

IV. Evaluation of Results

The chosen metric for training was the mAP, a standard in the evaluation of object detection models. This standard provides a comprehensive measure of precision across different IoU (Intersection over Union) thresholds. This metric ensures the model's effectiveness in accurately detecting objects, considering both the ability to correctly locate objects and to minimize false positives.

During the training, the model reached a stopping point at epoch 239 due to the EarlyStopping mechanism, which did not identify significant improvements in the last 50 epochs. The most important results were observed at epoch 189, where the model achieved an mAP of 0.453, evaluating the predictions with an IoU of 50% (mAP50). The model also achieved an mAP of 0.185 when considering the mAP 50-95. The F1-Score-Confidence curve illustrates the model's ability to balance precision and recall. This score is a composite metric that harmonizes the precision and recall of the model into a single indicator, ranging from 0 to 1, where 1 represents perfect performance. Confidence, in turn, reflects the probability attributed to detection by the model. By adjusting the confidence threshold, the F1 curve reveals the optimal point at which the model effectively balances precision and recall. In this case, it is observed that the model achieved a maximum F1 score of approximately 0.49 at a confidence point of 0.378. The results indicate an intrinsic challenge in discriminating between airstrip and background, showing a tendency for significantly higher false negatives than false positives. This tendency may be associated with the fact that some airstrips do not display features completely visible at this resolution level (10 meters per pixel). On the other hand, the relatively lower number of false positives reveals the network's capability to not confuse linear relief features with the targets of interest. Figure 5 presents the F1 curve in relation to confidence.

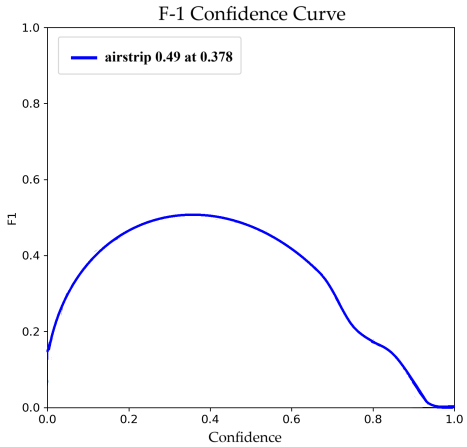

Fig. 5: Validation F-1 curve

At its best during the validation phase, the network correctly detected 133 of the 277 airstrips, while 144 were erroneously classified as false negatives (background). Additionally, in 50 instances, the model indicated an airstrip's presence where there was none (false positives). Table I presents the confusion matrix.

	Airstrip(True)	**Background(True)**
Predicted Airstrip	133	50
Predicted Background	144	–

TABLE I: Validation Confusion Matrix

A total of 117 airstrips were set aside for the generalization test. The targets were classified into three difficulty levels in terms of detection by visual inspection by a human analyst: easy, medium, and difficult. Considering these three levels, the number of targets was distributed as follows: 58 airstrips of difficult visibility, 19 medium level, and 40 easy.

V. Discussion

The initial analysis pertains to airstrips that are easily identifiable by human analysts. These targets are clearly visible in SAR images and do not have other structures nearby that could cause confusion. Out of 40 airstrips, 39 were correctly detected. Figure 6 presents two examples, a correct detection and the only easily detectable airstrip (with the portuguese term "pista" in the boudingbox) that was missed. This result indicates a high success rate of the network in detecting isolated airstrips or those that do not share significant similarities with adjacent structures. However, the failure to detect the other airstrip raises questions about the algorithm's limitations in the presence of complex urban elements. It is possible that the network was confused by overlapping patterns or the proximity of similar linear structures, such as streets. A hypothesis to improve the YOLOv8 network's accuracy in detecting airstrips in complex urban scenarios is to enhance training with a more diverse dataset, including numerous examples of airstrips adjacent to various urban forms. This could lead to better results in distinguishing airstrips from other linear structures not intended for this purpose.

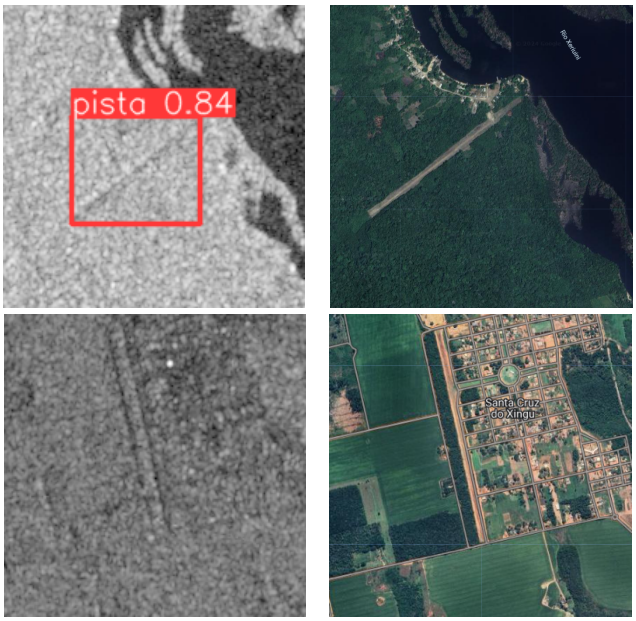

Fig. 6: Results for easy-to-see airstrips a) (lon: -61.9338 lat: -0.8771) b) (lon: - 52.3945, lat: -10.1587)

About the airstrips considered of medium difficulty, 6 of 19 were detected correctly. Of the medium-difficulty airstrips not detected, three cases corresponded to the presence of false positives. In these three cases, straight segments of streets or roads led to erroneous detection by the YOLOv8 network. Figure 7 exemplifies one of these cases. Of the 58 airstrips considered difficult, one target was detected correctly. Of the 57 cases not detected, 8 presented false positives. Of these, two cases had false positives generated due to the contrast presented by a cluster of trees and a deforested land. Figure 8 presents an example of each case, a correct detection and a wrong detection caused by a cluster of trees. The performance observed in medium and high-difficulty cases, while revealing certain challenges, underscores the potential for significant improvement with a more comprehensive training dataset. By utilizing a more robust and diverse database, encompassing a broader range of airstrips in various environments and conditions, the network's ability to generalize and accurately detect targets could be substantially enhanced. This approach would enable YOLOv8 to learn from a wider array of features and scenarios, thereby increasing its efficacy. Specifically, incorporating additional examples of airstrips in proximity to other linear infrastructures would refine the network's capability to distinguish airstrips from similar structures, leading to more reliable and accurate detections in complex settings.

VI. Concluding Remarks

This study demonstrated the capability of the YOLOv8 network in identifying unauthorized airstrips through SAR images in the Amazon region, utilizing a database composed of 646 images for training, 277 for validation, and 117 for generalization tests. The model's effectiveness in detecting the

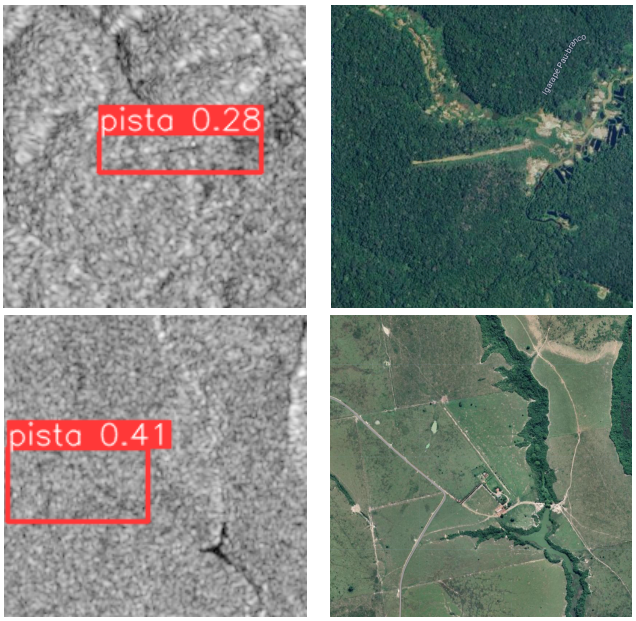

Fig. 7: Results for airstrips of medium difficulty to see a)(lon:-57.5819, lat: -6.8842) b) (lon: -48.4788, lat: -4.2370)

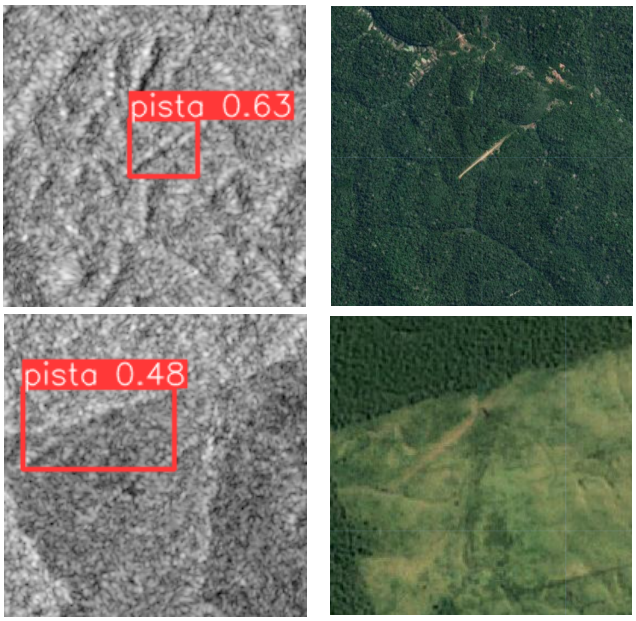

Fig. 8: Results for hard-to-see airstrips. a)(lon:-56.6452, lat:-5.5831) b)(lon: -49.7681 lat:-7.5793)

airstrips points to the possibility of using computer vision for monitoring vast forest areas. The precision achieved, despite facing difficulties distinguishing between actual airstrips and other natural linear features, shows a promising adaptability of the model for complex and densely vegetated scenarios like the Amazon. The importance of continuing the development and refinement of image processing and machine learning techniques is emphasized. Enhancing these technologies is crucial for increasing accuracy and reducing rates of false positives and negatives, thereby facilitating their practical application on a large scale.

A. Acknowledgments

The authors of this article would like to thank the Brazilian Air Force (FAB), the Postgraduate Program in Operational Applications (PPGAO), the Electronic Warfare Laboratory (LabGE-ITA), for their support during this work. This study was financed in part by the Brazilian agency Financier of Studies and Projects (FINEP) under Grant 01.22.0581.00.

REFERENCES

[1] IBGE. (2024) Amazônia legal. Accessed : 14 april 2024. [Online]. Available: https://www.ibge.gov.br/geociencias/cartas-e-mapas/mapas-regionais/15819-amazonia-legal.html

[2] I. S. (ISA), "Yanomami under attack: Illegal mining on yanomami indigenous land," Instituto Socioambiental, Technical Report, 2022, available online: https://acervo.socioambiental.org/acervo/documentos/yanomami-under-attack-illegal-mining-yanomami-indigenous-land (Accessed: 14 April 2024).

[3] A. Y. Hutukara and A. W. Ye'kwana, "Scars in the forest: evolution of illegal mining in the yanomami indigenous land in 2020," in *Monitoring System of Illegal Mining in the Yanomami Indigenous Land*, I. S. (ISA), Ed. Instituto Socioambiental (ISA), 2021, pp. 1–50.

[4] T. Kuck, P. Silva Filho, E. Sano, P. Bispo, E. Shiguemori, and R. Dalagnol, "Change detection of selective logging in the brazilian amazon using x-band SAR data and pre-trained convolutional neural networks," *Remote Sens.*, vol. 13, p. 4944, 2021, accessed: 12 april 2024. [Online]. Available: https://doi.org/10.3390/rs13234944

[5] L. Pang, B. Li, F. Zhang, X. Meng, and L. Zhang, "A lightweight yolov5-mne algorithm for SAR ship detection," *Sensors*, vol. 22, p. 7088, 2022.

[6] H. C. T. De Oliveira, "Detecção automática de pequenas embarcações em imagens SAR da constelação iceye com uso de uma rede neural convolucional," São José dos Campos, p. 84, 2021.

[7] T. Kuck, E. Sano, P. Bispo, E. Shiguemori, P. Silva Filho, and E. Matricardi, "A comparative assessment of machine-learning techniques for forest degradation caused by selective logging in an amazon region using multitemporal x-band SAR images," *Remote Sens.*, vol. 13, p. 3341, 2021, accessed: 12 abril 2024. [Online]. Available: https://doi.org/10.3390/rs13173341

[8] A. Moreira, P. Prats-Iraola, M. Younis, G. Krieger, I. Hajnsek, and K. Papathanassiou, "A tutorial on synthetic aperture radar," *IEEE Geoscience and Remote Sensing Magazine*, vol. March 2013, pp. 6–43, accessed: 20 maio 2024. [Online]. Available: https://doi.org/10.1109/MGRS.2013.2248301

[9] F. Meyer, "Spaceborne synthetic aperture radar: Principles, data access, and basic processing techniques," in *The SAR Handbook*. Fairbanks, AK, USA: University of Alaska Fairbanks, 2021, pp. 1–53.

[10] T. Zhang *et al.*, "SAR ship detection dataset (SSDD): Official release and comprehensive data analysis," *Remote Sens.*, vol. 13, p. 3690, 2021.

[11] S. Lei *et al.*, "Srsdd-v1.0: A high-resolution SAR rotation ship detection dataset," *Remote Sens.*, vol. 13, p. 5104, 2021.

[12] G. Suresh, C. Melsheimer, J.-H. Körber, and G. Bohrmann, "Automatic estimation of oil seep locations in synthetic aperture radar images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, pp. 4218–4229, 2015.

[13] K. Agarwal, M. Sanyo, S. Bakshi, V. M, J. J, and D. S, "Performance analysis of yolov7 and yolov8 models for drone detection," in *2023 International Conference on Network, Multimedia and Information Technology (NMITCON)*, 2023, pp. 1–6.

[14] Ultralytics. (2024) Yolo licenses: How is ultralytics yolo licensed. Accessed: 11 de abril de 2024. [Online]. Available: https://docs.ultralytics.com/#yolo-licenses-how-is-ultralytics-yolo-licensed

[15] Roboflow. (2024) What's new in yolov8. Accessed: 12 abril 2024. [Online]. Available: https://blog.roboflow.com/whats-new-in-yolov8/

[16] A. Elhanashi, P. Dini, S. Saponara, and Q. Zheng, "Telestroke: real-time stroke detection with federated learning and yolov8 on edge devices," *Journal of Real-Time Image Processing*, vol. 21, no. 121, pp. 1–16, 2024.

[17] Ultralytics. (2024) Yolov8: Métricas de desempenho. Accessed: 11 de abril de 2024. [Online]. Available: https://docs.ultralytics.com/pt/models/yolov8/#performance-metrics