



MSU Graduate Theses


Summer 2024

Road Extraction on Remote Sensing Imagery: Historical Mapping of the Brazilian Amazon

Jonas Paiva Botelho Jr
Missouri State University, Jonas031@live.missouristate.edu

As with any intellectual project, the content and views expressed in this thesis may be considered objectionable by some readers. However, this student-scholar's work has been judged to have academic value by the student's thesis committee members trained in the discipline. The content and views expressed in this thesis are those of the student-scholar and are not endorsed by Missouri State University, its Graduate College, or its employees.

Follow this and additional works at: <https://bearworks.missouristate.edu/theses>

 Part of the [Artificial Intelligence and Robotics Commons](#), [Data Science Commons](#), [Environmental Indicators and Impact Assessment Commons](#), [Geographic Information Sciences Commons](#), [Remote Sensing Commons](#), and the [Spatial Science Commons](#)

Recommended Citation

Botelho, Jonas Paiva Jr, "Road Extraction on Remote Sensing Imagery: Historical Mapping of the Brazilian Amazon" (2024). *MSU Graduate Theses*. 3984.
<https://bearworks.missouristate.edu/theses/3984>

This article or document was made available through BearWorks, the institutional repository of Missouri State University. The work contained in it may be protected by copyright and require permission of the copyright holder for reuse or redistribution.

For more information, please contact bearworks@missouristate.edu.

**ROAD EXTRACTION ON REMOTE SENSING IMAGERY: HISTORICAL MAPPING OF
THE BRAZILIAN AMAZON**

A Master's Thesis

Presented to

The Graduate College of
Missouri State University

In Partial Fulfillment

Of the Requirements for the Degree
Master of Science, Computer Science

By

Jonas Paiva Botelho Junnior

August 2024

Copyright 2024 by Jonas Paiva Botelho Junnior

ROAD EXTRACTION ON REMOTE SENSING IMAGERY: HISTORICAL MAPPING OF THE BRAZILIAN AMAZON

Computer Science

Missouri State University, August 2024

Master of Science

Jonas Paiva Botelho Junnior

ABSTRACT

This work proposes an artificial intelligence model based on U-Net architecture to map road networks in the Brazilian Amazon. Over the years, the Amazon region has been heavily exploited, leading to increased deforestation rates, contributing to CO₂ emissions, amplifying global warming, and causing a disturbance in local fauna and flora. The expansion into the forest by illegal miners, loggers, and land grabbers can be tracked down by the construction of roads, which we can refer to as the arteries of deforestation. Previous works on the matter proposed algorithms that use high-resolution imagery to map roads precisely. However, this work approach goes a step further and proposes the usage of medium-resolution imagery from Landsat satellite to map the roads of the Amazon, spatially and temporally. By taking advantage of years of data acquired by the Landsat satellite family, I was able to create a historical map of roads in the Amazon from 1985 to 2020, which provides information on the process of expansion of roads into the Amazon region. The model was trained on Landsat 8 imagery using more than 3 million kilometers of roads as its dataset, reaching 0.6577, 0.62371, and 0.6401 for precision, recall, and f1-score metrics. These results are important for predictive models of high-risk areas of deforestation, allowing for more precise estimation of roadless forests, having been made publicly available.

KEYWORDS: artificial Intelligence, remote Sensing, road network, deep learning, cloud-based platforms

**ROAD EXTRACTION ON REMOTE SENSING IMAGERY: HISTORICAL MAPPING OF
THE BRAZILIAN AMAZON**

By

Jonas Paiva Botelho Junnior

A Master's Thesis
Submitted to the Graduate College
Of Missouri State University
In Partial Fulfillment of the Requirements
For the Degree of Master of Science, Computer Science

August 2024

Approved:

Lloyd Smith, Ph.D., Thesis Committee Chair

Tayo Obafemi-Ajayi, Ph.D., Committee Member

Siming Liu, Ph.D., Committee Member

Jun Luo, Ph.D., Committee Member

Carlos Souza Jr, Ph.D., Committee Member

Julie Masterson, Ph.D., Dean of the Graduate College

In the interest of academic freedom and the principle of free speech, approval of this thesis indicates the format is acceptable and meets the academic criteria for the discipline as determined by the faculty that constitute the thesis committee. The content and views expressed in this thesis are those of the student-scholar and are not endorsed by Missouri State University, its Graduate College, or its employees.

ACKNOWLEDGEMENTS

I would like to thank my family for the support through all these years. It has not been easy, but it has been rewarding. All those video calls, the words of encouragement, weekly updates, and reminders of my responsibilities have helped me during the course of my graduate studies.

I would also like to thank my advisor and the committee for all the support, instruction, and patience in explaining and leading me in the right direction. I will always be grateful for all the work that we have accomplished together, not only in terms of the thesis but all the classes, assignments, and tests. To all my professors that helped me to acquire new knowledge and guided me into the final steps of my master's degree.

Last, but not least, I would like to thank my wife, that has been with me through every step, every doubt, every moment of sadness and relief. All the struggle and happiness. Thank you for helping me achieve one more milestone in this crazy journey called life.

I dedicate this thesis to Monica, Jonas, Bete, Deise, and Ana.

TABLE OF CONTENTS

Introduction	10
Literature Review	13
Deep Learning in Image Segmentation	13
Road Extraction Approaches	13
Road Extraction on Moderate-Resolution Imagery	14
Data pre-processing and post-processing	16
Model Tuning and Accuracy Assessment	17
Methodology	19
Study Area	20
Amazon Road Dataset 2020 (ARD)	21
Landsat Composites	21
Sample Processing	23
Post-Processing	25
U-Net Model	26
Model Assessment	29
Results	31
Metrics	31
Visualization	34
Large-scale inference and reference comparison	36
Brazilian Legal Amazon 2020 Landsat Road Network Map	36
Landsat model vs. Sentinel model	38
Time-Series data	42
Road development	43
Roads x Deforestation x Forest	44
Future Work	48
Conclusion	50
Reference	51

LIST OF TABLES

Table 1. Missing Tiles in Landsat Mosaic by Year.	23
Table 2. Performance metrics for pre and post-processing model's inference over testset.	34
Table 3. Timeseries road's statistics by year.	46

LIST OF FIGURES

Figure 1. Workflow describes data preprocessing, model training, and post-processing.	19
Figure 2. Map of the Brazilian Legal Amazon (study area).	20
Figure 3. Visual representation of image RGB composite after pre-processing. (a) displays Landsat 8 2020 composite. Visualization of each band that composes pre-processed Landsat images: (b) Shortwave Infrared 1 Band (Swir 1), (c) Near-Infrared (NIR), and (d) Red.	22
Figure 4. Batch processing and sample extraction locations.	24
Figure 5. Proposed neural network architecture.	27
Figure 6. Loss function minimization graph for training and validation sets.	32
Figure 7. Performance metrics development during model training. (a) Precision metric; (b) Recall metric; (c) F1 score metric.	33
Figure 8. Visual examples of the model's results on the test-set: (a) Input image, (b) reference data, (c) model's raw output, and (d) model's output after raster post-processing	35
Figure 9. Brazilian Legal Amazon's 2020 road network map produced by the proposed model.	36
Figure 10. Three examples of false positive residues. The (I) represents the false positives generated at the border of the charts. Images (II) showcase roads wrongly mapped in areas of drought. (III) false positives mapped at odd cloud formations.	37
Figure 11. Map displaying the sub-area representing the Brazilian Amazon Biome borders.	39
Figure 12. Visual representation of road detection between the Sentinel-2 and Landsat models for RFAs and RDAs.	40
Figure 13. Comparison between satellites' road network map.	41

Figure 14. Visual representation of road development in the municipality of Novo Progresso, along BR-163. (a) Showcase the status of the area in 1985 while (b) displays its configuration in 2020.	43
Figure 15. Road network growth over time.	44
Figure 16. Graph encompassing total roads, RDAs, and RFAs data growth in the time series in comparison to deforestation growth.	45
Figure 17. Representation of roads and deforestation growth over time.	47
Figure 18. Representation of road network growth from (a) 1985 to (b) 2020 around the Indigenous Land Mãe Maria.	49

1. Introduction

In recent years, several developments have occurred in computer vision, with modern artificial intelligence architectures greatly influencing various aspects of our daily lives, including data generation, utilization, and storage. In remote sensing, it has enabled researchers, governments, and entrepreneurs to assess project needs, address societal issues, and monitor environmental concerns [1]. The use of artificial intelligence has also improved the mapping and monitoring of one important variable: roads. Roads are one of the most critical aspects when considering any form of transportation, being the main pathways for people, goods, and services [2,3]. However, their use also poses negative consequences, particularly regarding environmental preservation [4]. Over the years, roads have penetrated previously untouched forested areas, most notably in the Amazon Rainforest, where opening pathways through the forest for illegal logging and mining has significantly increased deforestation rates [5,6]. To mitigate and reduce this alarming trend, researchers and government agencies have increasingly turned to deep-learning solutions [7,8] for detecting and proactively managing road construction.

Monitoring the Amazon Rainforest presents a unique challenge due to its vast proportions. Until recently, analysts have manually mapped roads in the region by analyzing remote sensing imagery. Not only was it a time-consuming task, but it was also prone to human bias. In [9], the authors devised a protocol to minimize errors by enhancing the linear characteristics of roads, enabling analysts to identify their structure and reduce the misclassification of rivers, property lines, power lines, and bodies of water. The protocol uses image bands from Landsat TM and ETM+ sensors [10],

namely Bands 3 (Red), 4 (NIR), and 5 (SWIR1), which facilitated visualization of exposed soil and differentiation from other linear features. Using this protocol, researchers and analysts could gather valuable information on the regional dynamics and store data that would subsequently be utilized to automate the mapping process.

Building upon this work, Botelho *et al* [8] proposed implementing a U-Net model to speed up the mapping process. The authors opted to apply the protocol to Sentinel-2 imagery (10m resolution) to obtain high-quality samples for model training. Their proposed model achieved a precision of 72% and successfully identified 3 million kilometers of roads across 5 million square kilometers of land, enabling faster and more up-to-date monitoring of the Amazon rainforest. The model was further utilized in the Previsia project [11] to incorporate road information into a statistical model of deforestation risk, facilitating the assessment of the impact of roads on deforestation dynamics and recent developments. However, the model was trained on Sentinel-2 imagery, which is restricted by the satellite's resolution and time constraints, thereby limiting the analysis to very recent data (2016 to the present).

Nevertheless, key questions remain unanswered regarding the historical evolution of road dynamics, particularly considering that the expansion of roads in the Brazilian Amazon gained momentum relatively recently, in the 1960-70s, with the Trans-Amazon highway construction plan [12]. How does the expansion of roads happen? Was it an organized process, or did the construction of new pathways happen haphazardly? What drove roads in a certain direction, and what can we observe regarding environmental preservation? The data presented in this thesis can provide insight into answering these questions.

Obtaining historical road data for the region is challenging, as it is either scarce or incomplete. The goal of the work presented here is to apply road extraction techniques developed for Sentinel 2 (10-meter) imagery to the medium resolution (30-meter) imagery provided by the Landsat program [10]. Such a model would enable the study of road expansion in the Amazon from 1985 to the present day. To that end, this study proposes an approach that utilizes deep learning techniques, recent road data, and Landsat imagery to train a model capable of detecting roads in lower-resolution data spanning over 30 years and applying time-series inference over 5-year snapshots to account for meaningful changes. The proposed approach leverages the data resources offered by Google Earth Engine [13] and Google Colaboratory [14] to acquire data, pre-process, and train an inference model capable of identifying roads throughout the Amazon rainforest. The subsequent sections describe the step-by-step process employed to obtain the data, pre-process it, feed it to the neural network, and utilize it to generate a historical road vector map.

2. Literature Review

2.1. Deep Learning in Image Segmentation

The field of image segmentation has grown significantly throughout the years, going from methods such as threshold [15] and clustering [16] to more complex structures like Generative Adversarial Networks (GANs) [17], autoencoders [18,19], and more recently, multi-model architecture that allows for the generating of masks from user-generated prompts [20]. These algorithms aim to create a pixel-wise classification map of an image. Autoencoders, in particular, are the most popular network structures used for image segmentation, given their feature extraction performance [19]. It reconstructs the image's structure from the extracted features by encoding the input features into a feature space and decoding it to predict desired classes, maintaining the same format between input and output.

Ronneberger, *et al* [21] proposed a variation of autoencoders named U-Net that extracted semantic information through an encoder path and precisely mapped pixel classes through a symmetrical decoder path. It was awarded the 2015 ESBI cell tracking challenge in the field of medical image segmentation and was deemed a state-of-the-art model for image segmentation. From there, variation of its structure expanded to other disciplines like remote sensing, where building and road extraction algorithms [22–24] took advantage of its ability to learn an expressive number of features from very few examples.

2.2. Road Extraction Approaches

The extraction of roads from images has been a longstanding and relevant topic in the literature. Numerous works addressing this issue can be traced back several

decades [1] and are of significant interest. The approaches proposed to tackle this problem encompass various techniques, from manual labeling to classical and modern machine learning. Earlier methods involved knowledge-based algorithms and mathematical solutions for classification that focused on spectral and contextual information present in the images [1].

In recent years, deep learning architectures have revolutionized image segmentation. A recent contribution in this field was made by the novel proposal in [21], which established the U-Net auto-encoder network as the state-of-the-art approach, delivering highly accurate results for image segmentation. Building upon this foundation, subsequent works, such as [22,24], have explored various modifications to the U-Net architecture to address specific challenges in road segmentation. For instance, recurrent units were incorporated to facilitate information propagation within the networks, and adjustments were made to the dimensionality of input and output data.

Additionally, some approaches have relied on a combination of models. For instance, the UU-Net [25] model proposed utilizing diverse data formats during training. This approach improved feature extraction capabilities and yielded promising results in road segmentation. Furthermore, the proposed approach used two distinct data resources, taking advantage of Sentinel 1 and Sentinel 2 imagery. Combining both databases allowed for a cleaner visualization of roads using optical and radar information.

2.3. Road Extraction on Moderate-Resolution Imagery

When employing deep learning algorithms for road extraction, it is essential to consider the quality of the input data, including both the labeling accuracy and the

resolution of the images used to train the model. The resolution of the images significantly affects the model's ability to perceive and identify roads accurately. While there have been few attempts to develop models capable of identifying roads in medium and high-resolution images, some notable works have addressed this challenge.

For instance, in [7], Gao, et al. utilized aerial datasets such as the Massachusetts Roads Dataset, which offers a resolution of less than 1 meter. On the other hand, other works [8,25,26] have attempted to bridge the gap between moderate and high-resolution imagery, ranging from 30 meters to 1 meter [27]. The issue is not only related to the data quality but also its availability. Higher-resolution imagery is often limited to private data providers and comes at a high cost, while moderate and low-resolution imagery are more accessible for research. Moreover, lower-resolution imagery usually offers broader coverage, increasing its appeal for road extraction tasks in locations far from city centers.

Against this backdrop, Kearney, *et al.* [26] employed RapidEye satellite data in mapping roads in rural areas of Canada, augmenting the process with on-ground GPS trackers to establish more robust labels, achieving high recall and precision, 0.89 and 0.87, respectively. Lin, *et al.* [25], utilized both Sentinel-1 and Sentinel-2 data to train a unified UU-Net model, leveraging the combined information to improve road extraction accuracy and mitigate issues related to cloud noise, reaching a F1-Score value of 0.56. Botelho, *et al.* [8], the proposed approach combined remote sensing techniques to enhance linear features in Sentinel-2 images and introduced a variant of the U-Net model to identify roads in the Brazilian Amazon, generating a model with 0.71 of precision and recall of 0.65.

It is worth noting that early attempts to automate road mapping relied on Landsat imagery, which offered a resolution of 30 meters, utilizing knowledge-based systems and image processing techniques [9]. However, these earlier approaches did not fully exploit the potential of deep learning and were limited by the coarser resolution of the data.

2.4. Data pre-processing and post-processing

As noted earlier, much of the existing work in road detection has predominantly relied on using very high-resolution imagery. Several well-known datasets, such as the Massachusetts Roads Dataset [28], OpenStreetMap [29], and GF-2 Road Dataset [13], have served as valuable resources for these studies. In each case, pre-processing and post-processing techniques were employed to enhance the accuracy of the results.

The pre-processing stage can be categorized into two main aspects: image data pre-processing and label data pre-processing. Image data pre-processing involved subdividing the input image into smaller samples for feeding into the training process. Also, image normalization techniques were commonly employed to ensure consistent pixel values across the dataset. For instance, Chen, *et al.* [31], used a 37949×35341 image of a central part of New York to train their neural network. They cut the original image into 1368 images of 1000×1000 pixels to ease the training process and divided them into three sets of data (e.g., training, validation, and test).

In their post-processing phase, Kearney, *et al.* [26], took measures to enhance the accuracy of the road detection results. This was achieved by connecting segments and removing misclassified data. Botelho, *et al.* [8] adopted a different approach by generating a binary mask through thresholding to distinguish between road and non-

road areas. The predicted patches were combined to form a comprehensive study area map. Subsequently, the raster information was converted to vector format to extract statistical data, such as length, quantity, and density of roads. By employing these pre-processing and post-processing techniques, researchers have improved the accuracy and reliability of road detection algorithms, enabling more effective road mapping and analysis.

2.5. Model Tuning and Accuracy Assessment

Hyperparameter tuning plays a crucial role in AI training, as it involves adjusting the model's variables to optimize accuracy measurements. The tuning process involves modifying parameters such as learning rate, loss functions, activation functions, optimizers, dropout rates, and model size (length and depth). By fine-tuning these characteristics, models can achieve varying levels of performance.

Different loss and activation functions have been explored to address specific challenges. Zhang, *et al.* [24], used mean squared error to minimize loss, while the Rectified Linear Unit (ReLU) was applied to the residual layer as the activation function to filter ≤ 0 values. Havaei, *et al.* [32] employed the generalized dice loss to handle unbalanced data and a maxout non-linearity function to model features. Moreover, in the original U-Net publication [21], Ronnenberg introduced an asymmetrical input and output correlation, whereas Botelho, *et al.* [8], made modifications to preserve edge pixels and ensure input and output shapes were the same.

Once these models are trained, several well-known accuracy metrics are employed to evaluate their performance, particularly for unbalanced data such as roads. Sun, *et al.* [33] employed precision, recall, and F1-Score, along with other metrics such as mIoU

and IoU, to assess the quality of their results. Detailed explanations of these metrics will be provided in the subsequent chapters.

3. Methodology

The methodology for training an AI model capable of identifying roads in the Amazon region consists of three steps, outlined in Figure 1. The first step involves acquiring and preparing the training data for input into the neural network training process. The second step focuses on the actual training of the model and the analysis of its accuracy parameters. Lastly, the third step encompasses the post-processing phase, which involves cleaning and filling any gaps in the output. Each of these steps will be described in detail in the following sections.

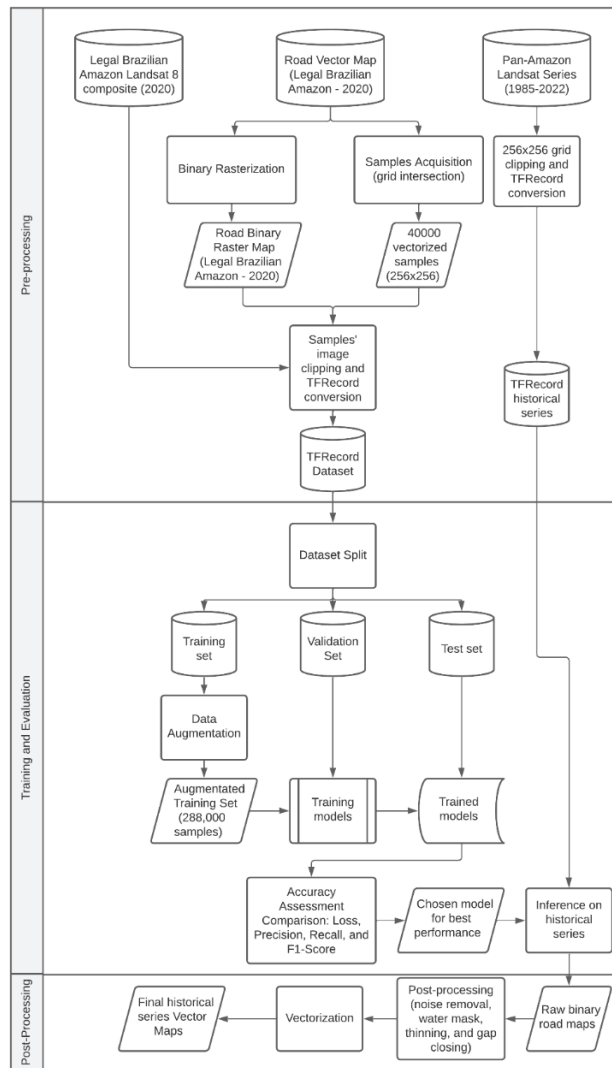


Figure 1. Workflow describes data preprocessing, model training, and post-processing.

3.1. Study Area

The study area encompasses the entire Legal Brazilian Amazon [34], which includes the nine states that contain parts of the Amazon biome and serve as the location for collecting training samples for the proposed AI model, as illustrated in Figure 2. The vast area covers approximately 5 million square kilometers and contains different road formations and Land use scenarios, such as pastureland, mining, and logging. The diversity facilitates the training process in effectively generalizing the characteristics distinguishing roads from non-roads. Moreover, the dataset generated by Botelho, *et al.* [8] was used to divide the region into smaller samples to train the neural network. These training samples are rectangular structures of 256×256 pixels in size that define the geolocation where the data pre-process will clip the Landsat RGB composite and the rasterized binary road data.

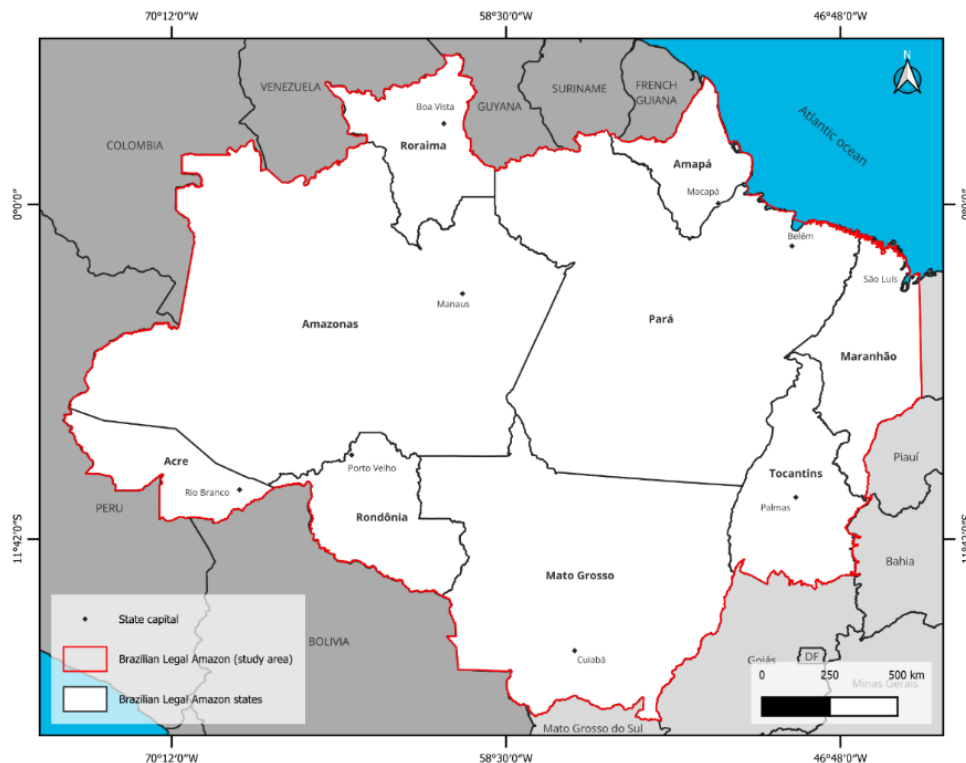


Figure 2. Map of the Brazilian Legal Amazon (study area).

3.2. Amazon Road Dataset 2020 (ARD)

The training process of the proposed AI model utilized the dataset generated automatically by the U-Net model presented in [8] as reference data. This dataset encompasses a comprehensive collection of roads spanning 3 million kilometers within the Brazilian Legal Amazon region, specifically from 2020. The original data format is a vector file in .shp format, which underwent two fundamental processes: sample location and binary mask conversion. The sample location process involved the selection of relevant tiles from a grid that covers the entire study area. This process used the vector file to identify samples intersecting with road vectors, as illustrated in Figure 3.

Simultaneously, the vector data was converted into a raster file with a resolution of 30 meters to align with Landsat's specifications. Prior to rasterization, the vector data was buffered to twice its original size to better represent the roads and their surrounding context. The dataset is heavily unbalanced due to the nature of roads as a rare class in the Brazilian Amazon's overall context, with road pixels representing only 258 million pixels (13%). In contrast, non-road pixels represent 1.8 billion pixels (88%). To deal with the imbalance, checkpoints, performance metrics, and calculation functions were set to ensure minimal negative impact on the learning process.

3.3. Landsat Composites

This work utilized two satellites from the Landsat family to create the yearly mosaics: Landsat 5 and 8, each covering 1985-2010 and 2013-present, respectively. I applied a pre-processing routine that included the following steps: selecting specific bands to enhance the linear features of roads, defining the temporal constraint to retrieve the desired images, adjusting the brightness and gain values, and applying standard

deviation stretch to sharpen the imagery. Depending on the satellite that generated the imagery, the parameters of the steps were slightly modified to achieve similar results.

For the training dataset, the image mosaic for the year 2020 was generated by filtering image tiles corresponding to August and October, thus avoiding cloud interference [35]. Additionally, the images were filtered based on their cloud cover score, which was set to $\leq 30\%$, reducing noise in the dataset. The mosaics were then combined by calculating the median value for each pixel from the specified timeframe. A histogram stretch applying 2 standard deviations was employed to mitigate the bright spikes present in the images. Furthermore, gain values were adjusted for each band (0.08, 0.06, and 0.2), and a gamma value of 0.75 was applied to enhance visualization of road features. The pre-processing result is represented by Figure 3. Similarly, the image mosaics for 1985, 1990, 1995, 2000, 2005, 2010, and 2015 underwent pre-processing to conform to the training data format.

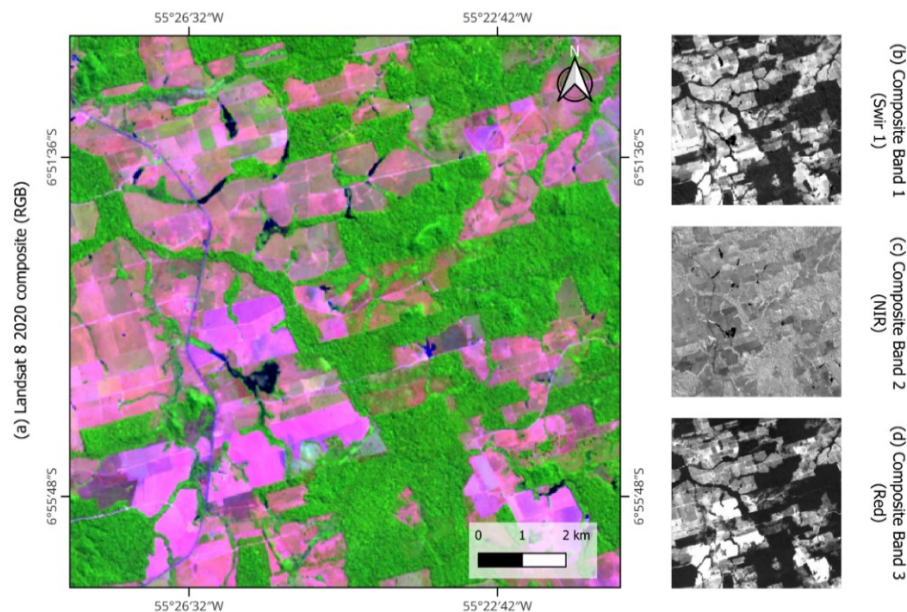


Figure 3. Visual representation of image RGB composite after pre-processing. (a) displays Landsat 8 2020 composite. Visualization of each band that composes pre-processed Landsat images: (b) Shortwave Infrared 1 Band (Swir 1), (c) Near-Infrared (NIR), and (d) Red.

While accessing images from the Google Earth Engine database was straightforward, yearly mosaics often lacked completeness, particularly in older images. These gaps resulted from the cloud cover filter, compounded by the absence of satellite imagery in those regions. The most significant gaps were observed in the northern part of the Amazon, characterized by prevalent year-round cloud cover. The years 1985 and 2010 had the most missing tiles, representing 8.76% and 7.79% of the entire study area, respectively. Several alternatives were available to address this issue, including incorporating pixels from different years into the mosaic process and generating multi-year mosaics. However, I kept data from a single year in the Amazon yearly mosaics to preserve year-by-year information. Table 1 delineates the Landsat scenes absent from the inventory and their corresponding missing areas.

Table 1. Missing Tiles in Landsat Mosaic by Year.

Year	N° of missing tiles	Missing area (km ²)	Missing Area (%)
1985	14	439,448	8.76
1990	7	177,738	3.54
1995	10	254,841	5.08
2000	5	158,700	3.16
2005	1	31,510	0.62
2010	13	391,219	7.79
2015	0	0	0.0
2020	0	0	0.0

3.4. Sample Processing

Training samples involved systematically selecting relevant areas characterized by roads within the specific regions. These samples were standardized to 256 × 256 pixels and extracted based on a grid system encompassing the entirety of the Brazilian Legal Amazon. Approximately 40,000 samples were selected based on the presence of roads from the reference data, covering an estimated 2.3 million square kilometers. To

facilitate computational operations, the samples were extracted in batches following the structure of the International Millionth Map of the World sheets [36], as visually represented in Figure 4.

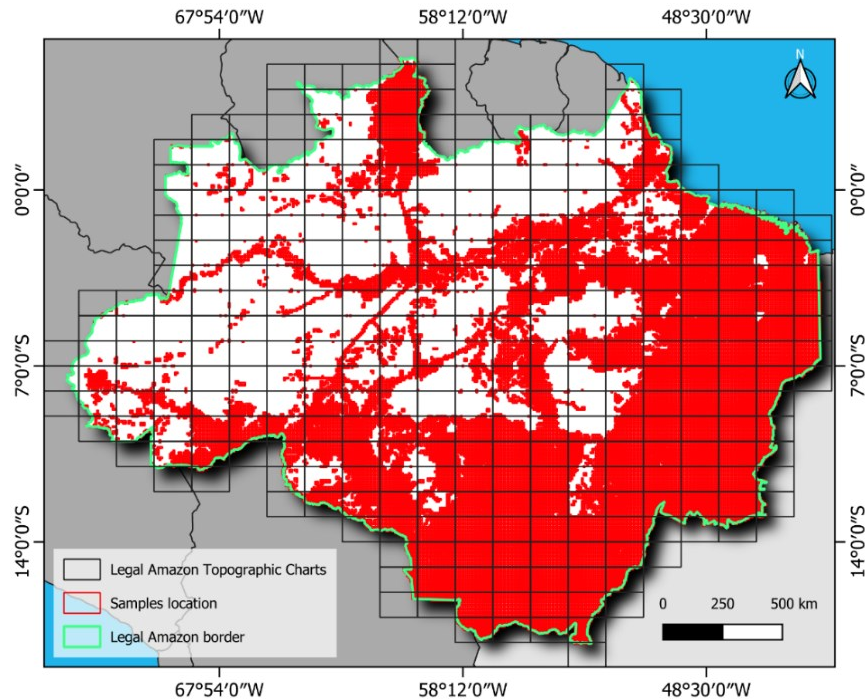


Figure 4. Batch processing and sample extraction locations.

Each iteration of the process utilized the geometry of the sheets to create a corresponding grid. Subsequently, the grid was filtered based on its intersection with road vectors, resulting in the selection of sample locations. Once the locations were determined, the binary and Landsat images were cropped and combined to construct the dataset structure, consisting of Input/Image data and Label Data. These structures were converted from .tiff format to numpy arrays to facilitate data management. Furthermore, the numpy arrays were organized by bands and stored in TFRecord files to streamline the data ingestion process within the TensorFlow [37] framework. In total, 334 TFRecord files were created based on the International Millionth Map of the World Sheets that covers the training area.

The files were randomly split into training, validation, and test datasets. This allocation designated 80% of the data for training, with the remaining 20% evenly distributed between validation and test sets. Data augmentation techniques were applied to the training set to augment sample diversity and enhance the model's generalization. Each training sample underwent rotations of 90, 180, and 270 degrees and horizontal and vertical flips, increasing the sample count from 32,000 to 160,000.

3.5. Post-Processing

The model generates a $256 \times 256 \times 1$ image array containing class probability values for each pixel. An empirical decision was made to set the confidence threshold at 0.5, effectively filtering out pixel values with less than a 50% probability of representing a road. In addition to the thresholding process, several post-processing steps were implemented to enhance the model's results and to convert the output from an image format into georeferenced raster and vector data.

Following the elimination of low-probability pixels, the array values are georeferenced using the coordinates of the original input image. Georeferencing the image is essential for pinpointing the precise location of roads and applying geoprocessing tools for further data analysis. Additionally, a water mask [38] is applied over the image to remove residual artifacts resembling linear structures that may occur between bodies of water and land, such as sand bodies.

The post-processing algorithm closes gaps between segments separated by as few as 2 pixels before thinning the structures to a 1-pixel length. This gap-closing operation involves applying a geometric structure buffer, rectangles of size 2×1 , to search for connections between nearby segments before undergoing thinning. Once the segments

are interconnected and thinned, the post-processing algorithm leverages the Python library GRASS GIS [39] to execute geoprocessing operations, including Vectorization, Generalization, Border Cleaning, and Density Cleaning.

The vectorization process plays a crucial role in transforming the output data from raster to vector format, leveraging the location of single-pixel-width segments. These segments adhere to a pixel-squared format while constituting the road vector representation. Subsequently, this representation undergoes a generalization process to smooth out lines, ensuring a more accurate portrayal of roads. The vector data enables the analysis of segment lengths, facilitating mapping the road network's extent.

Despite the effectiveness of vectorization, misclassified border pixels, often resulting from tilted scenes and distorted images, posed challenges. To address this issue, a border-cleaning process was implemented. This involved utilizing auxiliary vector files representing these borders to identify and remove road segments aligned with them. Once the borders were cleaned, the entire dataset was divided into patches of 10 kilometers. This segmentation allowed for the examination of road density, aiming to eliminate small false positives introduced by factors such as clouds, unmapped rivers, and bodies of sand.

3.6. U-Net Model

I opted to use a U-Net architecture as it is considered state-of-the-art for image segmentation and because of its past performances in segmenting roads from medium-resolution remote sensing imagery (10 meters). The modified U-Net architecture proposed for road mapping in the Amazon region was developed based on previous work of Botelho, et al. [8]. It has undergone further refinements to adapt to differences in

the input data. This modified architecture has increased depth to enhance the feature extraction process by allowing the model to acquire more high-level data due to Landsat's low resolution. The architecture comprises five layers of encoder and decoder blocks, with a total of ≈ 3.1 million parameters, organized into four primary block steps: Encoder Convolutional Block, Bottleneck Convolutional Block, Decoder Convolutional Block, and Dropout Convolutional Block, illustrated in Figure 5.



Figure 5. Proposed neural network architecture.

The Encoder Convolutional Block plays a pivotal role in extracting high and low-level features by applying a sequence of 2-dimensional convolutional layers coupled with batch normalization. The Maxpooling layer is also employed in this block. Given the high data imbalance in the dataset, the Leaky ReLU activation function ($\alpha = 0.3$) [40] was chosen to mitigate the issue of dying neurons. Notably, the results from each block are preserved before the max-pooling operation to facilitate data reconstruction by the subsequent decoder convolutional blocks. In contrast, the Decoder Convolutional Blocks utilize transposed convolution operations to fill gaps in the image and reconstruct its information while adapting their weights based on reference data. The Bottleneck and Dropout Convolutional Blocks are intermediate steps for extracting low-level features and introducing variation and regularization into the training process.

Additionally, the model employed the Nadam optimization algorithm with a learning rate of 0.001 and the soft dice loss function [41], as described in equation 1, to adjust its weights during backpropagation. This function computes the Dice Coefficient by measuring the overlap between the model's output probabilities and the reference data. The numerator represents the doubled count of positive activations, while the denominator denotes the total number of activations. The coefficient is then subtracted from 1 to facilitate minimization of the results, effectively mitigating the impact of class imbalances.

$$(1) \text{ Dice Coefficient} = \frac{|2 * \sum y_{true} * y_{pred}|}{|\sum y_{true}^2 + \sum y_{pred}^2 + 1|}$$

$$\text{Soft Dice} = 1 - \text{Dice Coefficient}$$

Regarding training parameters, the dataset was input into the model in batches of 32 samples over 200 epochs. Additionally, checkpoints, early stopping, and learning

rate reduction mechanisms were implemented to ensure optimal convergence and efficient utilization of computing resources. These three mechanisms are triggered based on the stability of the validation loss function, which is evaluated after a series of non-improvement iterations.

The implementation of the model can be accessed through the following link <https://colab.research.google.com/drive/19FQCRnzJezmqY5S9tgVcvdT3tSZLL9Ho?usp=sharing>

3.7. Model Assessment

Additionally, alongside the loss function value, the training process computed other performance metrics to gauge the learning progress of the neural network. Precision, recall, and F1-score metrics were selected for their reliability in assessing the quality of the model's outcomes. These metrics leverage fundamental variables such as true positives, false positives, and false negatives to quantify the accuracy of predictions. Precision assesses the model's accuracy by considering true and false positives, as defined by Eq 2. It evaluates the model's ability to predict positive outcomes.

$$(2) \textit{Precision} = \frac{\textit{True Positives}}{\textit{True Positives} + \textit{False Positives}}$$

Furthermore, Recall evaluates the model's completeness using true positives and false negatives. Including false negatives allows for assessing the model's accuracy in predicting positive outcomes.

$$(3) \textit{Recall} = \frac{\textit{True Positives}}{\textit{True Positives} + \textit{False Negatives}}$$

Lastly, the F1-score calculates the harmonic mean between the precision and recall rates, offering a balanced evaluation of the model's performance.

$$(4) F_1 = \frac{2 * Precision * Recall}{Precision + Recall}$$

In addition to the statistical metrics, I analyzed the data from two additional perspectives. First, the model's performance in large-scale inference allows for assessing the model's ability to identify roads in the entire study area. Areas not used during the model's training might present different contexts that can potentially generate more data. The large-scale inference was compared with Sentinel's model map of 2020 to compare data statistical information (e.g., length and presence). Second, I applied the model over the time series and analyzed how the roads changed over time against one another. These analyses are presented in chapters five and six, respectively.

4. Results

The model underwent training in the Google Colab environment, utilizing a configured NVIDIA V100 GPU and a system equipped with 50GB of RAM. The training process was completed in approximately ~8 hours, yielding the trained model's archive and a log file encompassing crucial performance metrics. These metrics include the loss function value, precision, recall, and F1-score for the training and validation sets.

Despite being set to train for 200 epochs, the model demonstrated early convergence and negligible improvement after reaching a plateau. Consequently, the early stopping parameter intervened at epoch 102. Upon concluding the training process, the checkpoint parameter ensured that the model's weights were set to those of the best epoch, determined based on the validation loss value—precisely, epoch 77.

Furthermore, as previously detailed, a distinct test set was employed to assess the model's generalization capability, evaluating its proficiency in discerning the presence of roads in previously unseen image data. Comprehensive details of these results are explained in the following sections.

4.1. Metrics

Three key aspects were scrutinized to assess the model's effectiveness: the convergence of the training and validation sets and the performance on the test set. Throughout training and validation, the model's primary indicator of convergence was the minimization of the soft dice loss function, as shown in Figure 6. This metric, starting at 25.11 and 2.42 for training and validation, respectively, steadily decreased until reaching its lowest point at epoch 77—0.3321 for training and 0.3286 for validation. Concurrently, the loss function value during the test evaluation returned as 0.3070.

These values, all derived from epoch 77, were selected as the best epoch based on the parameters considering the lowest validation loss function value. The loss function value was chosen due to its non-interfering nature during training and its representation of the complete model without applying the dropout layer.




Figure 6. Loss function minimization graph for training and validation sets.

In addition to the loss function, precision, recall, and F1-score metrics were employed to validate the completeness and correctness of the results by comparing the model's output to the reference data pixel-wise. The precision metric reached a maximum of 0.6341 during training and 0.6408 during validation, while recall values were 0.5867 and 0.5934, respectively. The F1-score, representing a harmonic balance between precision and recall, achieved values of 0.6085 for training and 0.6156 for validation. These metrics indicate that compared to the reference data, the model completed a match of approximately 60% during training and validation, as shown in Figure 7. Moreover, the test set demonstrated model performance in new data,

returning metrics values of 0.6577 for precision, 0.6237 for recall, and 0.6401 for the F1-score.

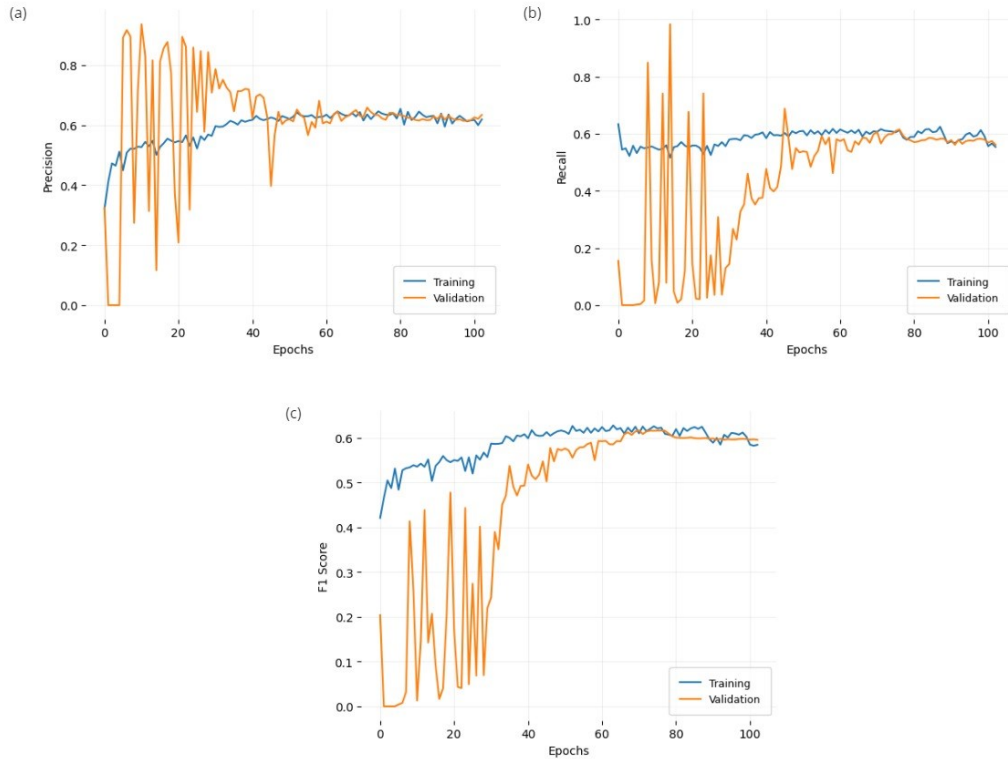


Figure 7. Performance metrics development during model training. (a) Precision metric; (b) Recall metric; (c) F1 score metric.

The metrics mentioned above were calculated directly from the model's output and represent the model's performance without any external interference. However, as mentioned previously, a post-process is applied to the results to improve segment connectivity and clean possible false positives. To grasp the performance of the entire workflow, the metrics were also calculated for the post-processed data over the test set, which indicates a slightly lower precision (0.6388), recall (0.5700), and f1-score (0.6024), described in Table 2. Despite lower metrics, the results still hold good quality, affected by small distortion caused by dilation and inflation processes applied to connect segments, as shown in Figure 8.

Table 2. Performance metrics for pre and post-processing model's inference over testset.

	Precision	Recall	F1-Score
Testset raw output	0.6577	0.6237	0.6401
Testset post-processed output	0.6388	0.5700	0.6024

4.2. Visualization

To evaluate the quality of the results, various examples generated by the model are analyzed, as in Figure 8. The examples illustrate different road contexts and formations, organized in rows, with varying inference stages displayed in columns. Column A shows the input data, Column B depicts the reference data, and Columns C and D show the model's output and the results after the post-processing steps.

Comparing the reference data to the model's output reveals that the reference data, sourced from higher-resolution images, contains more roads. However, the proposed model maintains prominent roads and visible structures in the Landsat imagery. More minor roads, particularly those near exposed soil, are more challenging to identify, resulting in fragmented segments in the output. Roads surrounded by forests are easier to identify due to their distinct linear features, as observed in row 4. Additionally, their organized context makes geometric road networks more easily identifiable. Overall, the model's outputs provide meaningful results, capturing most road network information despite missing some segments due to resolution limitations and mixed class signals.

The post-processing steps applied to the results enhance road sharpness, close gaps between fragmented segments, and reduce noise. However, these steps also slightly modify the structure of some segments, mainly due to inflation and dilation processes. Some squared-shaped road patches become more spherical, and bordering

road segments may merge. Despite these changes, they maintain connectivity and effectively represent dynamics in the regions.




Figure 8. Visual examples of the model's results on the test-set: (a) Input image, (b) reference data, (c) model's raw output, and (d) model's output after raster post-processing

5. Large-scale inference and reference comparison

5.1. Brazilian Legal Amazon 2020 Landsat Road Network Map

The proposed model generated results for the entire study area, mapping its road networks in medium-resolution satellite imagery, facilitating broader spatial and temporal analysis of road dynamics. The entire Brazilian Legal Amazon region was utilized to validate this claim, employing the 2020 Landsat mosaic composite to produce a comprehensive road map. This process, conducted following the established protocol for acquiring training tiles, took approximately 8 hours to complete. The model mapped approximately 3 million roads across 5 million square kilometers, with a concentration primarily observed in the southern part of the study area, as depicted in Figure 9.

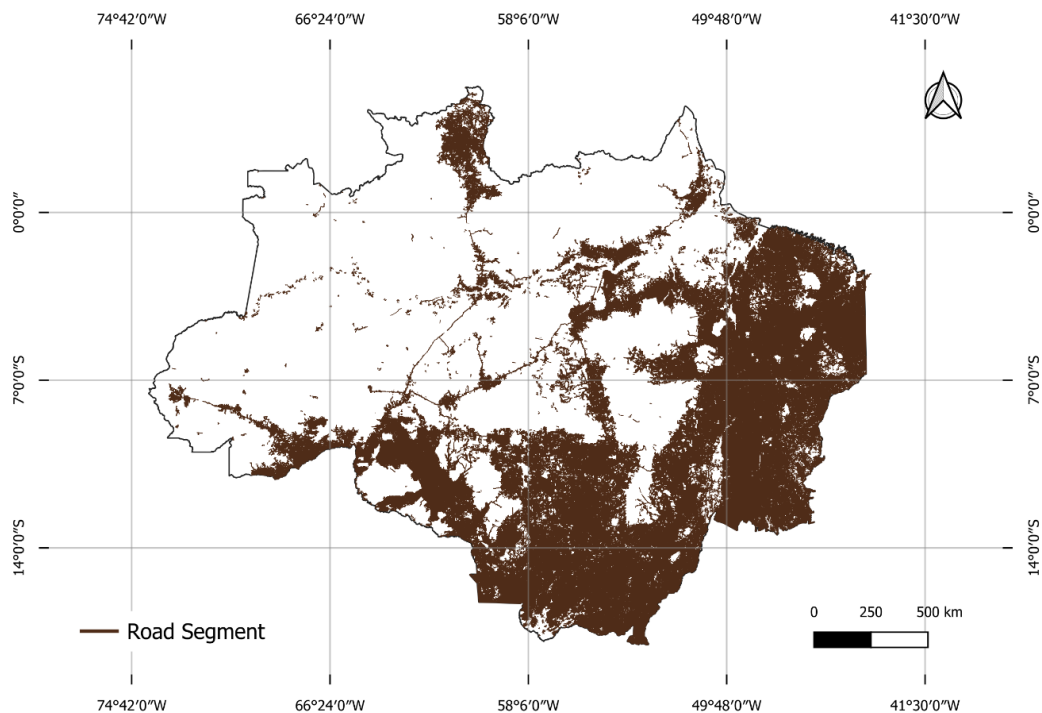


Figure 9. Brazilian Legal Amazon's 2020 road network map produced by the proposed model.

Two additional post-processing steps were implemented to refine the results. Firstly, a manual verification consisted of visually inspecting areas close to rivers, image borders, and areas with clouds to clean residual false positives, such as irregular cloud formations, drought areas, and exposed soil close to image borders, as shown in Figure 10. Once identified, those features were manually excluded from the overall road map, reducing the number of roads mapped to approximately 2.8 million kilometers.

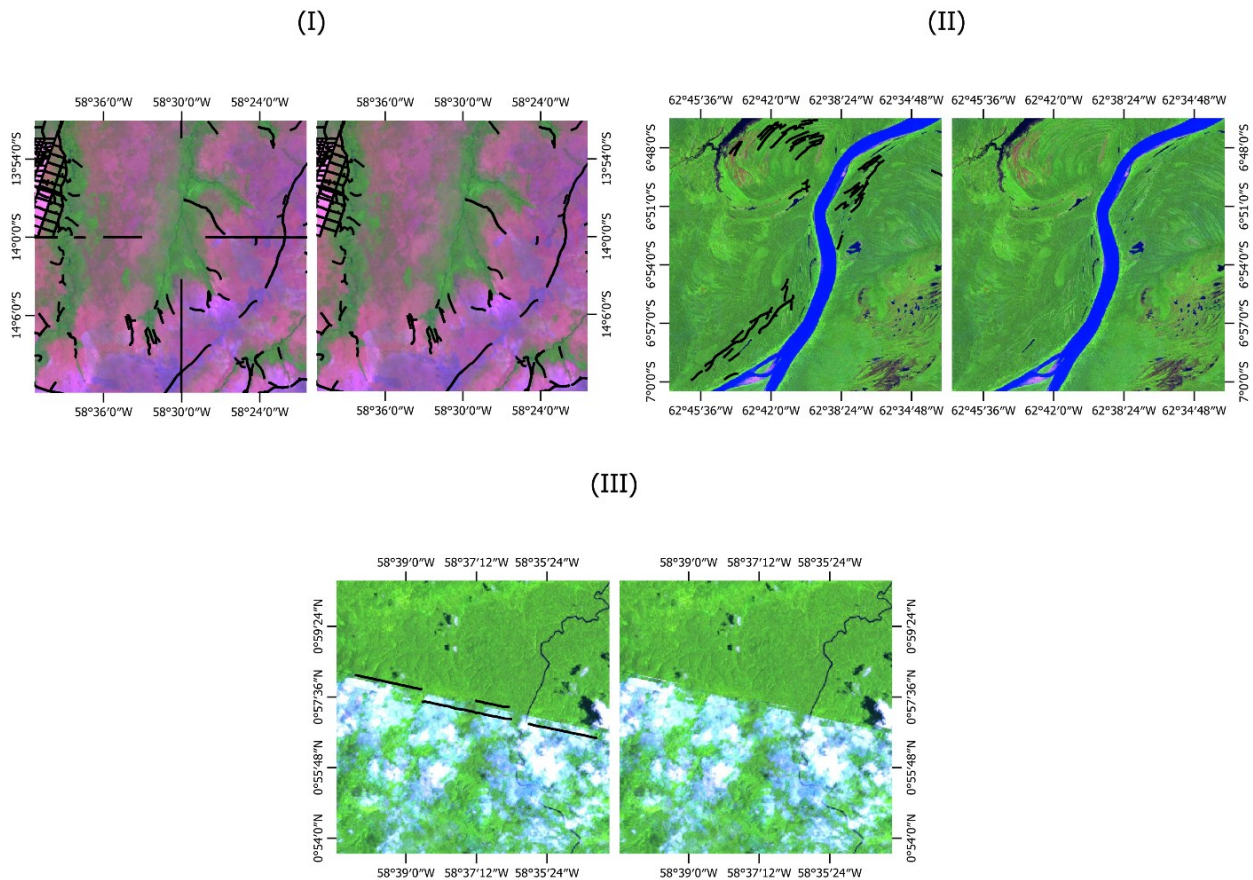


Figure 10. Three examples of false positive residues. The (I) represents the false positives generated at the border of the charts. Images (II) showcase roads wrongly mapped in areas of drought. (III) false positives mapped at odd cloud formations.

The second step involves consolidating connected segments in the .shp file to reduce complexity. Although the vectorization process effectively produced road line strings, it fragmented connected segments, increasing file size. The 2020 vectorized

output contained over 12 million features stored in a 1.7 GB .shp file, resulting in slow loading times in geospatial software like QGIS and Google Earth Engine. To address this issue, I employed the PostGIS database query function `st_clusterdbscan` [42] to identify connected segments within a pixel distance of 30 meters or ≈ 0.000269 degrees. Once clustered, segments were merged using the `st_union` function, resulting in a ≈ 1 GB file with approximately 164 thousand rows, simplifying the dataset.

5.2. Landsat model vs. Sentinel model

Benchmarking against the ARD map of the Brazilian Legal Amazon was crucial to assess the model's performance comprehensively. Despite being trained using ARD, the proposed model's inference across the entire study area could unveil new insights, especially in unmapped regions. Given the challenge of comparing road-by-road similarities, a spatial distribution analysis was conducted instead to ascertain whether the proposed model's data aligned spatially with the reference data, indicating its usability.

To address potential oversights in hotspots, deforestation data from [43] was utilized to categorize roads into two categories: Roads in Deforested Areas (RDA) and Roads in Forested Areas (RFA). This categorization enabled separate analysis of high and low-density regions, enhancing the accuracy of the spatial distribution assessment. Moreover, the spatial distribution analysis was confined to the limits of the Brazilian Amazon Biome to maintain geographical consistency and relevance, as shown in Figure 11.

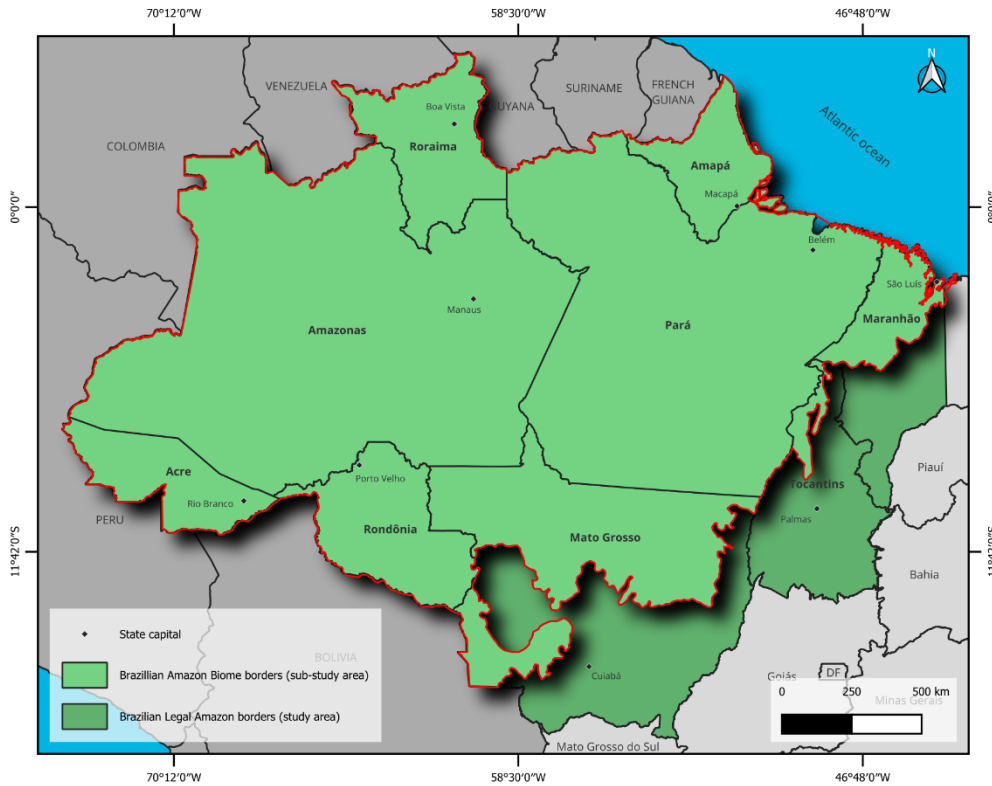


Figure 11. Map displaying the sub-area representing the Brazilian Amazon Biome borders.

The analysis focused on a subarea spanning 4.26 million km², revealing 1.73 million km of roads identified by the proposed model and 2.15 million km by the ARD. For Roads in Deforested Areas (RDAs), the Landsat analysis returned 1.4 million kilometers of roads, representing an 11% reduction compared to the ARD, distributed across 63.7 Mha of deforested areas. Differences were expected due to the lower spatial resolution of Landsat relative to Sentinel. Both results exhibited similar road structures and pathways, as illustrated in Figure 12.

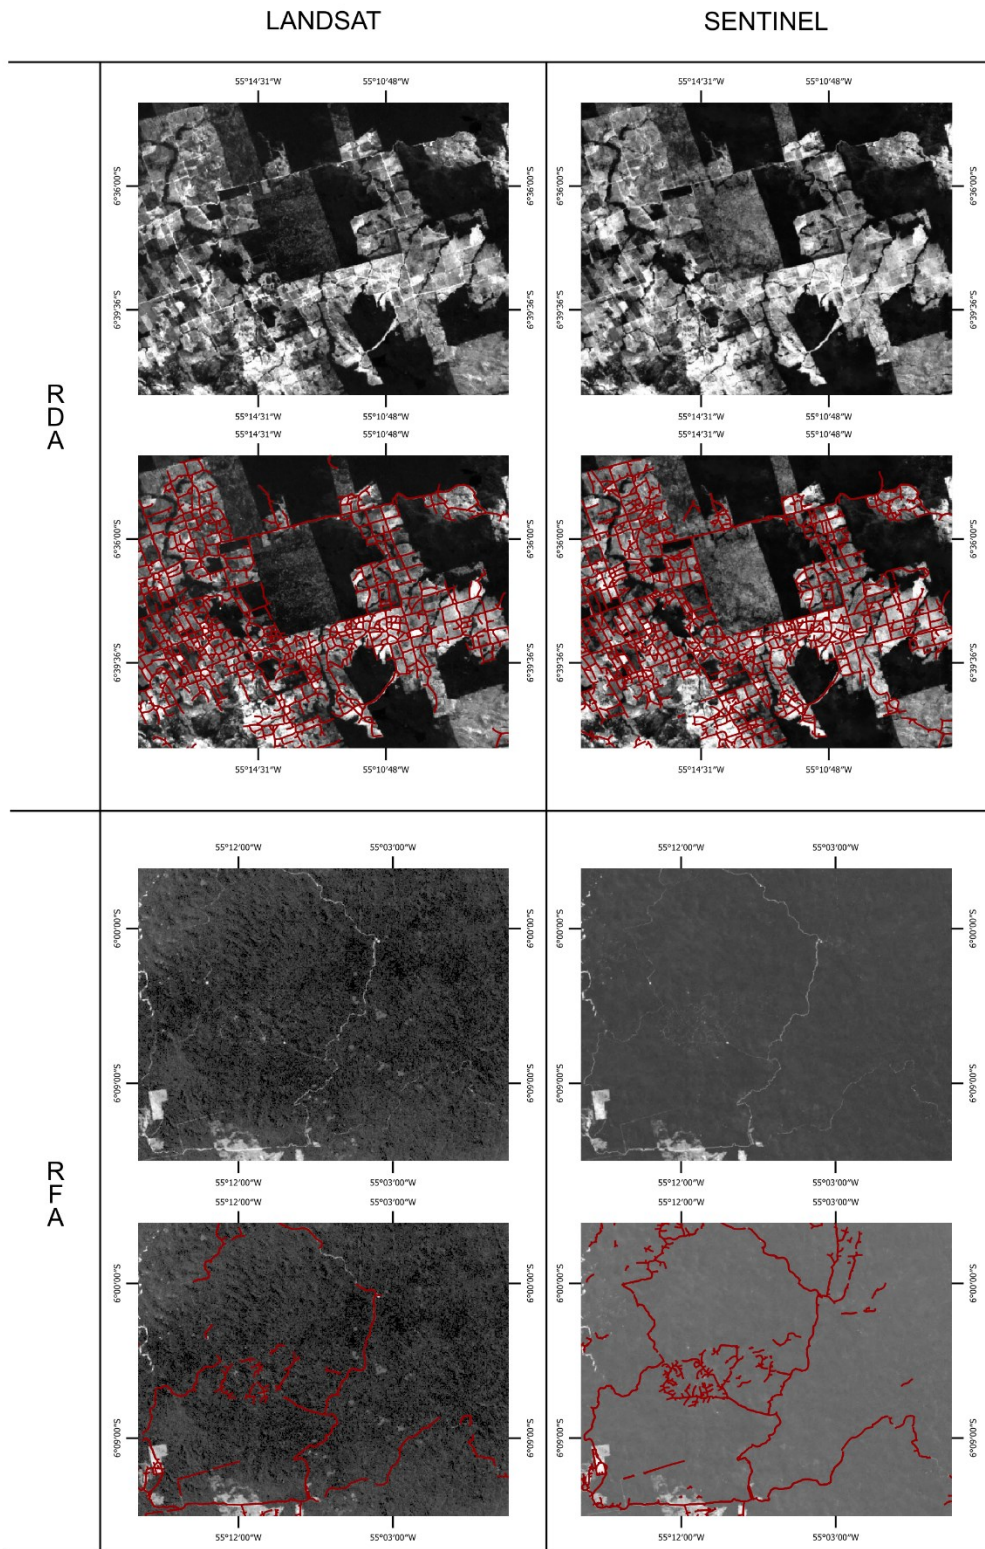


Figure 12. Visual representation of road detection between the Sentinel-2 and Landsat models for RFAs and RDAs.

The disparities primarily arise from smaller segments, which are particularly challenging to discern in Landsat imagery, such as those within areas of exposed soil and dense vegetation cover. Roads in Forested Areas (RFAs) are notably affected by the latter, often located in regions with minimal forest loss. As a result, the proposed model identified only 337 thousand kilometers of RFAs, contrasting with the 571 thousand kilometers detected by the ARD. Figure 13 illustrates the comparative proportions of both results, indicating that while the Landsat model can adequately map general roads, it struggles to detect roads visible only in higher resolutions.

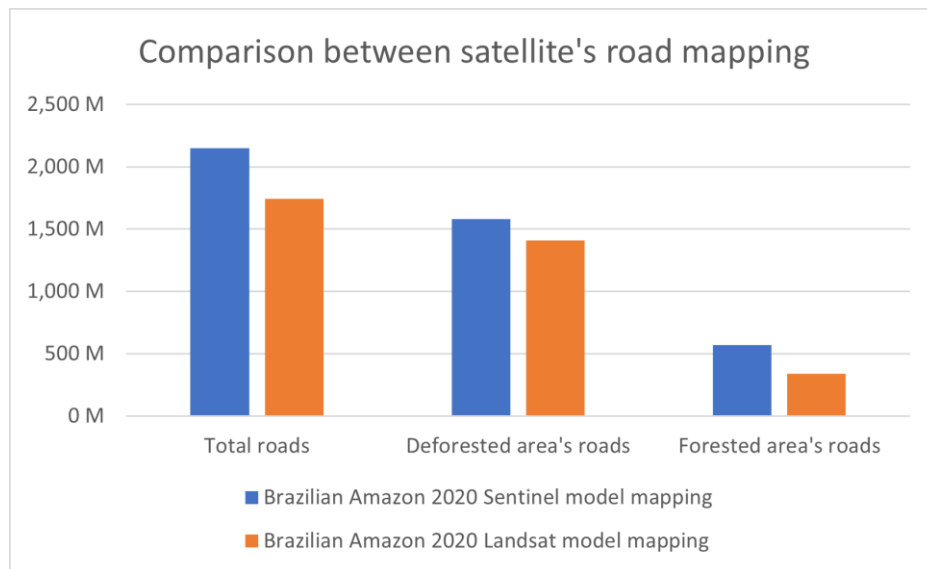


Figure 13. Comparison between satellites' road network map.

6. Time-Series data

Proposing a model capable of road inference in Landsat 8 imagery enables analysis of images from other satellites within the Landsat family, exploiting their spectral wavelength and resolution similarities. With over 30 years of Landsat satellite data available, the model facilitates mapping Brazil's road evolution. The Landsat imagery data spanning a 5-year cycle was utilized to assess its capabilities, focusing on road dynamics in the study area. Additionally, road network growth was analyzed concerning deforestation and forested area reduction, topics discussed further ahead. Figure 14 showcases the changes that happened in the municipality of Novo Progresso between 1985 and 2020.




Figure 14. Visual representation of road development in the municipality of Novo Progresso, along BR-163. (a) Showcase the status of the area in 1985 while (b) displays its configuration in 2020.

6.1. Road development

The model produced eight vector files spanning road network data from 1985 to 2020. Following the protocol outlined for the Landsat 8 2020 composite, the files were generated with a 5-year interval to highlight perceptible differences. Notably, Figure 14 demonstrates significant variations in mapped roads for each year, reflecting consistent growth. Results revealed an average 25% increase in the road network every five years,

with peaks observed during 2000-2005 and 2010-2015 – showing growth rates of 44% and 41%, respectively. However, the slow growth between 2005 and 2010 may be attributed to heightened cloud cover, reducing road mapping and causing discrepancies. Additionally, road expansion, albeit steady, shows regional disparities, with faster growth observed in the southern regions of the Brazilian legal Amazon, as depicted in Figure 15.




Figure 15. Road network growth over time.

6.2. Roads x Deforestation x Forest

Going further, I analyzed the time series results concerning deforestation. Historical data enables an examination of how the relationship between roads and deforestation has evolved. I conducted a comparative analysis between the results of the proposed Landsat model and those of [8] that used Sentinel-2. The road network was categorized into Roads in Deforested Areas (RDAs) and Roads in Forested Areas (RFAs). This

allowed for a year-by-year comparison to assess growth and correlation, detailed in Figure 16.

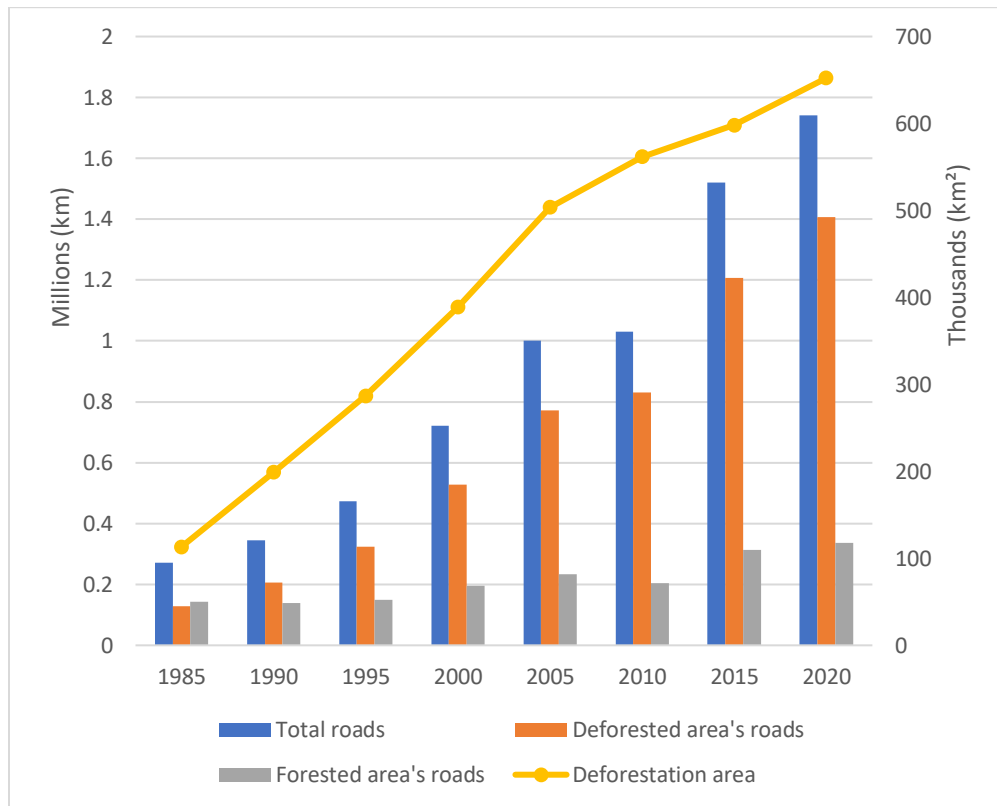


Figure 16. Graph encompassing total roads, RDAs, and RFAs data growth in the time series in comparison to deforestation growth.

The deforestation data indicates 63.7 Mha were deforested between 1985 and 2020, with a corresponding growth of 1.7 million kilometers in road data. Both experienced similar growth rates, with the road network expanding by an average of 31% over 35 years, paralleling the 30% increase in forest loss area, as illustrated in Graph 13. In 1985, approximately 1.28×10^5 kilometers of roads traversed 1.13×10^5 km² of deforested areas, while 1.43×10^5 km extended into the forest. Over 35 years, RDAs grew by 1095%, reaching over 1.4×10^6 km in 2020. Notably, RDAs exhibited two significant spikes: between 1985 and 1990, with a 61% increase from 1.28×10^5

km to 2.07×10^5 km, and between 1995 and 2000, with a 62% growth from 3.24×10^5 to 5.27×10^5 kilometers. The entire time-series data is described in Table 3.

Table 3. Timeseries road's statistics by year.

Years	Roads in Deforested Areas (RDAs)		Roads in Forested Areas (RFAs)		Total Roads	
	Length (10^6 x km)	Growth (%)	Length (10^6 x km)	Growth (%)	Length (10^6 x km)	Growth (%)
1985	0.12	-	0.14	-	0.27	-
1990	0.20	61.19	0.13	-2.95	0.34	27.33
1995	0.32	56.80	0.15	7.84	0.47	37.09
2000	0.52	62.40	0.19	30.27	0.72	52.23
2005	0.77	46.46	0.23	18.75	1.00	38.84
2010	0.83	7.66	0.20	-12.29	1.03	3.12
2015	1.20	45.14	0.31	53.46	1.52	46.76
2020	1.40	16.64	0.33	7.71	1.74	14.79

In contrast, RFAs provide discrete information on the state of roads within the forest each year, potentially transitioning into RDAs as surrounding land loses forest cover and is classified as deforested. Therefore, I compared RFAs based on their yearly extents rather than considering the data as incremental. This approach revealed fluctuations in RFAs, with reductions of -2% and -12% observed between 1985-1990 and 2005-2010, respectively, and from 2010 to 2015, the period witnessed a notable 53% increase in RFAs. However, missing Landsat scenes notably influenced this growth in 2010. Conversely, the second-highest growth rate occurred from 1990 to 1995, at 30%. Figure 17 visually represents the time-series of both deforestation and road networks, in which it is possible to notice that the older the deforestation is, the older the roads around it are. The same can be stated about newer deforested areas and roads.




Figure 17. Representation of roads and deforestation growth over time.

In conclusion, RDAs exhibited an average growth rate of 42%, while RFAs averaged approximately 14%. This suggests a more significant surge of roads in already deforested areas, indicating a more substantial influence on deforestation expansion into RDAs. However, it's important to note the temporal gap of 5 years in the analyzed data, limiting the ability to assess the impact of RFAs on deforestation directly. Further detailed analysis is needed to infer the specific effect of RFAs, which falls beyond the scope of this study.

7. Future Work

The development of road detection models for high-resolution imagery has accelerated data generation compared to manual mapping. This trend supports further study using high-resolution imagery and the creation of models focusing on broader coverage and higher temporal resolutions. Our Landsat Road detection model evolved from manual mapping to an automated AI detection model using medium-resolution imagery. This model enables retrospective analysis of roads in the Brazilian Legal Amazon over 35 years, revealing growth from 628 thousand to 2.89 million kilometers.

Benchmarking the model's performance from 1985 to 2015 relied on visual comparison with Landsat imagery composites. A direct comparison with the ARD wasn't feasible due to resolution differences, but an indirect comparison using deforestation data suggested comparable inference rates. The proposed AI model identified 1.74 million kilometers of roads in the Brazilian Amazon Biome, compared to 2.15 million by the ARD, primarily explained by resolution discrepancies. RDA represented 80% and 73% of the datasets, while RFA accounted for 19% and 25%.

The reliability of the data enables various research analyses, including deforestation evolution and human trafficking pathway changes. Historical datasets facilitate prediction systems' implementation, enhancing capabilities like the deforestation risk model [44]. Moreover, deforested areas have been growing rapidly throughout the years, except in protected lands, such as Indigenous Lands and Conservation Units, which play an essential role in environmental preservation. Visual analysis of road time-series data displays the effectiveness of these areas in reducing the expansion of networks and, consequently, deforestation, as shown in Figure 18. Further analysis can

be made to assess in detail how policy changes and monitoring of these lands interact with the progress of roads into forested territory. In addition, the correlation between roads and deforestation also opens the door to other subjects with overlapping issues, such as modern slavery [45], which necessitate further analysis to understand its impact and identify patterns for illegal site tracking.

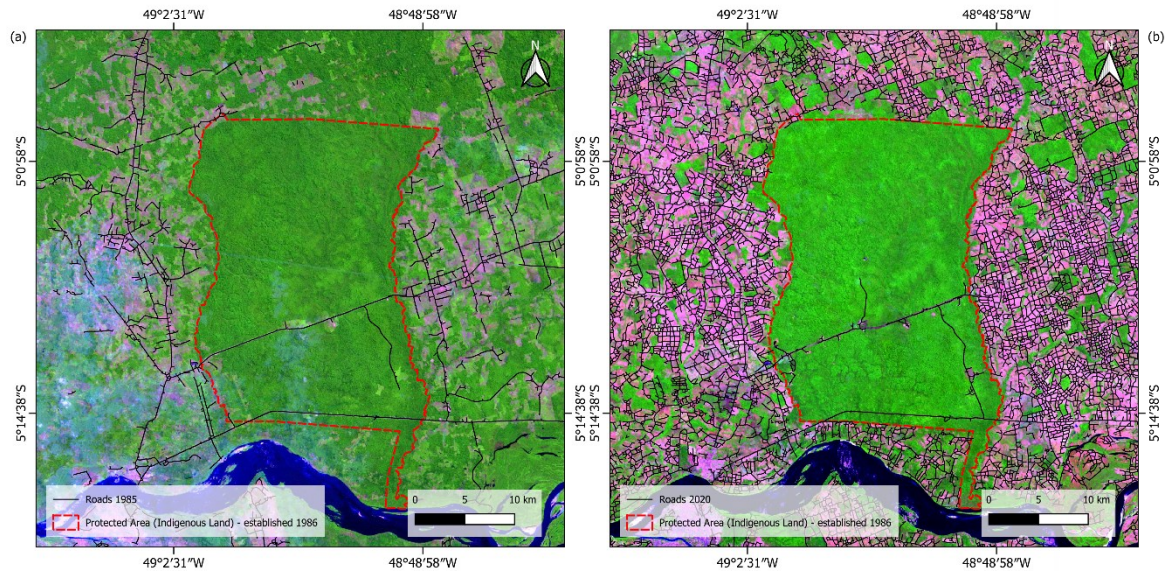


Figure 18. Representation of road network growth from (a) 1985 to (b) 2020 around the Indigenous Land Mãe Maria.

Improvements to the model input and post-processing are possible. Using multiyear mosaics instead of yearly ones could address missing data issues. However, this raises challenges in analyzing road growth and classifying representative data. Addressing topology arrangement and segment connectivity challenges requires manual adjustments, as automated solutions only partially resolve these issues.

8. Conclusion

The accumulation of decade-long manual mapping combined with AI algorithms enriched the overall data pool on roads, culminating in the Landsat Road detection model presented in this work. Also, by consolidating a pipeline that can consistently predict roads using new data, I created a time-series data for roads in the Brazilian Amazon, spanning over 30 years. The ever-evolving landscape and dynamics of the region, such as deforestation, fires, transportation of commodities, and trafficking, can be traced back in time and analyzed to answer questions regarding how we got to today's regional configuration. Over 30 years of road data can be used to improve predictive models to identify areas with higher deforestation risk, allowing for a more precise estimate of roadless forests. Furthermore, the data could be used to track the development of trafficking routes and their ramifications through time. Also, initial assessments of the correlation between deforestation and roads indicate a heavier influence in deforestation from roads in already deforested areas. Furthermore, improvements can still be made to ensure more complete results, targeting better topological representation, image composition, and segment connectivity. In addition, the model showcased its capability to generate high-quality results in low spatial and high temporal resolutions, which created a historical dataset for roads, opening the possibility for a diverse range of research and applications that can enrich how issues are tackled in the Amazon region.

9. Reference

1. Wang, W.; Yang, N.; Zhang, Y.; Wang, F.; Cao, T.; Eklund, P. A Review of Road Extraction from Remote Sensing Images. *Journal of Traffic and Transportation Engineering (English Edition)* **2016**, *3*, 271–282, doi:10.1016/J.JTTE.2016.05.005.
2. Milletari, F.; Navab, N.; Ahmadi, S.A. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In Proceedings of the Proceedings - 2016 4th International Conference on 3D Vision, 3DV 2016; 2016.
3. Xu, Y.; Xie, Z.; Feng, Y.; Chen, Z. Road Extraction from High-Resolution Remote Sensing Imagery Using Deep Learning. *Remote Sensing 2018, Vol. 10, Page 1461* **2018**, *10*, 1461, doi:10.3390/RS10091461.
4. Engert, J.E.; Campbell, M.J.; Cinner, J.E.; Ishida, Y.; Sloan, S.; Supriatna, J.; Alamgir, M.; Cislowski, J.; Laurance, W.F. Ghost Roads and the Destruction of Asia-Pacific Tropical Forests. *Nature* **2024**, *629*, 370–375, doi:10.1038/s41586-024-07303-5.
5. Perz, S.; Brilhante, S.; Brown, F.; Caldas, M.; Ikeda, S.; Mendoza, E.; Overdeest, C.; Reis, V.; Reyes, J.F.; Rojas, D.; et al. Road Building, Land Use and Climate Change: Prospects for Environmental Governance in the Amazon. *Philosophical Transactions of the Royal Society B: Biological Sciences* **2008**, *363*, 1889–1895, doi:10.1098/RSTB.2007.0017.
6. Southworth, J.; Marsik, M.; Qiu, Y.; Perz, S.; Cumming, G.; Stevens, F.; Rocha, K.; Duchelle, A.; Barnes, G. Roads as Drivers of Change: Trajectories across the Tri-National Frontier in MAP, the Southwestern Amazon. *Remote Sens (Basel)* **2011**, *3*, 1047–1066, doi:10.3390/rs3051047.
7. Gao, L.; Song, W.; Dai, J.; Chen, Y. Road Extraction from High-Resolution Remote Sensing Imagery Using Refined Deep Residual Convolutional Neural Network. *Remote Sens (Basel)* **2019**, *11*, 1–16, doi:10.3390/rs11050552.
8. Botelho, J.; Costa, S.C.P.; Ribeiro, J.G.; Souza, C.M. Mapping Roads in the Brazilian Amazon with Artificial Intelligence and Sentinel-2. *Remote Sensing 2022, Vol. 14, Page 3625* **2022**, *14*, 3625, doi:10.3390/RS14153625.
9. Brandão, A.O.; Souza, C.M. Mapping Unofficial Roads with Landsat Images: A New Tool to Improve the Monitoring of the Brazilian Amazon Rainforest. *Int J Remote Sens* **2006**, doi:10.1080/01431160500353841.
10. Landsat Missions | U.S. Geological Survey Available online: <https://www.usgs.gov/landsat-missions> (accessed on 8 July 2024).
11. PrevisIA | Fazer Do Futuro Sustentável a Nossa Única Previsão. Available online: <https://previsia.org.br/> (accessed on 17 July 2024).

12. Walker, R.; Perz, S.; Arima, E.; Simmons, C. The Transamazon Highway: Past, Present, Future. *Engineering Earth* **2011**, 569–599, doi:10.1007/978-90-481-9920-4_33.
13. Gorelick, N.; Hancher, M.; Dixon, M.; Ilyushchenko, S.; Thau, D.; Moore, R. Google Earth Engine: Planetary-Scale Geospatial Analysis for Everyone. *Remote Sens Environ* **2017**, doi:10.1016/j.rse.2017.06.031.
14. Bisong, E.; Bisong, E. Google Colaboratory. In *Building Machine Learning and Deep Learning Models on Google Cloud Platform*; 2019.
15. Jardim, S.; António, J.; Mora, C. Image Thresholding Approaches for Medical Image Segmentation - Short Literature Review. *Procedia Comput Sci* **2023**, 219, 1485–1492, doi:10.1016/J.PROCS.2023.01.439.
16. Chen, Z.; Qi, Z.; Meng, F.; Cui, L.; Shi, Y. Image Segmentation via Improving Clustering Algorithms with Density and Distance. *Procedia Comput Sci* **2015**, 55, 1015–1022, doi:10.1016/J.PROCS.2015.07.096.
17. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Networks. *Commun ACM* **2020**, 63, 139–144, doi:10.1145/3422622.
18. Augustauskas, R.; Lipnickas, A. Improved Pixel-Level Pavement-Defect Segmentation Using a Deep Autoencoder. *Sensors* **2020**, Vol. 20, Page 2557 **2020**, 20, 2557, doi:10.3390/S20092557.
19. Minaee, S.; Boykov, Y.; Porikli, F.; Plaza, A.; Kehtarnavaz, N.; Terzopoulos, D. Image Segmentation Using Deep Learning: A Survey. *IEEE Trans Pattern Anal Mach Intell* **2022**, 44, 3523–3542, doi:10.1109/TPAMI.2021.3059968.
20. Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A.C.; Lo, W.Y.; et al. Segment Anything. *Proceedings of the IEEE International Conference on Computer Vision* **2023**, 3992–4003, doi:10.1109/ICCV51070.2023.00371.
21. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proceedings of the Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*; 2015.
22. Wang, H.; Miao, F. Building Extraction from Remote Sensing Images Using Deep Residual U-Net. *Eur J Remote Sens* **2022**, 55, 71–85, doi:10.1080/22797254.2021.2018944/FORMAT/EPUB.
23. Ji, S.; Wei, S.; Lu, M. Fully Convolutional Networks for Multisource Building Extraction from an Open Aerial and Satellite Imagery Data Set. *IEEE Transactions*

- on *Geoscience and Remote Sensing* **2019**, *57*, 574–586, doi:10.1109/TGRS.2018.2858817.
24. Zhang, Z.; Liu, Q.; Wang, Y. Road Extraction by Deep Residual U-Net. *IEEE Geoscience and Remote Sensing Letters* **2017**, *15*, 749–753, doi:10.1109/lgrs.2018.2802944.
 25. Lin, Y.; Wan, L.; Zhang, H.; Wei, S.; Ma, P.; Li, Y.; Zhao, Z. Leveraging Optical and SAR Data with a UU-Net for Large-Scale Road Extraction. *International Journal of Applied Earth Observation and Geoinformation* **2021**, *103*, 102498, doi:10.1016/j.jag.2021.102498.
 26. Kearney, S.P.; Coops, N.C.; Sethi, S.; Stenhouse, G.B. Maintaining Accurate, Current, Rural Road Network Data: An Extraction and Updating Routine Using RapidEye, Participatory GIS and Deep Learning. *International Journal of Applied Earth Observation and Geoinformation* **2020**, *87*, doi:10.1016/J.JAG.2019.102031.
 27. Friedl, M.A.; Woodcock, C.E.; Olofsson, P.; Zhu, Z.; Loveland, T.; Stanimirova, R.; Arevalo, P.; Bullock, E.; Hu, K.T.; Zhang, Y.; et al. Medium Spatial Resolution Mapping of Global Land Cover and Land Cover Change Across Multiple Decades From Landsat. *Frontiers in Remote Sensing* **2022**, *3*, 894571, doi:10.3389/FRSEN.2022.894571/BIBTEX.
 28. Bandara, W.G.C.; Valanarasu, J.M.J.; Patel, V.M. SPIN Road Mapper: Extracting Roads from Aerial Images via Spatial and Interaction Space Graph Reasoning for Autonomous Driving. *Proc IEEE Int Conf Robot Autom* **2022**, 343–350, doi:10.1109/ICRA46639.2022.9812134.
 29. Grinberger, A.Y.; Minghini, M.; Juhász, L.; Yeboah, G.; Mooney, P. OSM Science—The Academic Study of the OpenStreetMap Project, Data, Contributors, Community, and Applications. *ISPRS International Journal of Geo-Information* **2022**, *Vol. 11, Page 230* **2022**, *11*, 230, doi:10.3390/IJGI11040230.
 30. Segmentation on Gaofen Satellite Image Available online: <https://www.kaggle.com/datasets/yaroslavnaychuk/satelliteimagesegmentation> (accessed on 7 March 2024).
 31. Chen, Z.; Wang, C.; Li, J.; Xie, N.; Han, Y.; Du, J. Reconstruction Bias U-Net for Road Extraction from Optical Remote Sensing Images. *IEEE J Sel Top Appl Earth Obs Remote Sens* **2021**, *14*, 2284–2294, doi:10.1109/JSTARS.2021.3053603.
 32. Ranjbarzadeh, R.; Bagherian Kasgari, A.; Jafarzadeh Ghousechi, S.; Anari, S.; Naseri, M.; Bendeche, M. Brain Tumor Segmentation Based on Deep Learning and an Attention Mechanism Using MRI Multi-Modalities Brain Images. *Scientific Reports* **2021**, *11:1* **2021**, *11*, 1–17, doi:10.1038/s41598-021-90428-8.

33. Sun, Z.; Zhou, W.; Ding, C.; Xia, M. Multi-Resolution Transformer Network for Building and Road Segmentation of Remote Sensing Image. *ISPRS International Journal of Geo-Information* 2022, Vol. 11, Page 165 **2022**, 11, 165, doi:10.3390/IJGI11030165.
34. Legal Amazon | IBGE Available online: <https://www.ibge.gov.br/en/geosciences/maps/regional-maps/17927-legal-amazon.html?edicao=18047> (accessed on 30 April 2024).
35. Wang, J.; Chagnon, F.J.F.; Williams, E.R.; Betts, A.K.; Renno, N.O.; Machado, L.A.T.; Bisht, G.; Knox, R.; Bras, R.L. Impact of Deforestation in the Amazon Basin on Cloud Climatology. *Proc Natl Acad Sci U S A* **2009**, 106, 3670–3674, doi:10.1073/pnas.0810156106.
36. Fundação Instituto Brasileiro de Geografia e Estatística. Biblioteca Central.; Fundação Instituto Brasileiro de Geografia e Estatística. Gerência de Documentação e Biblioteca. Mapas e Outros Materiais Cartográficas Na Biblioteca Central Do IBGE. **1983**.
37. Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. TensorFlow: A System for Large-Scale Machine Learning. In Proceedings of the Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016; 2016.
38. Pekel, J.F.; Cottam, A.; Gorelick, N.; Belward, A.S. High-Resolution Mapping of Global Surface Water and Its Long-Term Changes. *Nature* 2016 540:7633 **2016**, 540, 418–422, doi:10.1038/nature20584.
39. Neteler, M.; Bowman, M.H.; Landa, M.; Metz, M. GRASS GIS: A Multi-Purpose Open Source GIS. *Environmental Modelling and Software* **2012**, 31, 124–130, doi:10.1016/j.envsoft.2011.11.014.
40. Maas, A.L.; Hannun, A.Y.; Ng, A.Y. Rectifier Nonlinearities Improve Neural Network Acoustic Models. in *ICML Workshop on Deep Learning for Audio, Speech and Language Processing* **2013**, 28.
41. Sudre, C.H.; Li, W.; Vercauteren, T.; Ourselin, S.; Jorge Cardoso, M. Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations. In Proceedings of the Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); 2017.
42. Ester, M.; Kriegel, H.-P.; Sander, J.; Xu, X. A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. **1996**.
43. Souza, C.M.; Oliveira, L.A.; de Souza Filho, J.S.; Ferreira, B.G.; Fonseca, A. V.; Siqueira, J. V. Landsat Sub-Pixel Land Cover Dynamics in the Brazilian Amazon.

Frontiers in Forests and Global Change **2023**, *6*,
doi:10.3389/FFGC.2023.1294552.

44. Sales, M.; de Bruin, S.; Herold, M.; Kyriakidis, P.; Souza, C. A Spatiotemporal Geostatistical Hurdle Model Approach for Short-Term Deforestation Prediction. *Spat Stat* **2017**, *21*, 304–318, doi:10.1016/J.SPASTA.2017.06.003.
45. Stanford Researchers Investigate Human Trafficking Alongside Brazilian Partners in the Amazon Rainforest | FSI Available online:
<https://healthpolicy.fsi.stanford.edu/news/stanford-researchers-investigate-human-trafficking-alongside-brazilian-partners-amazon> (accessed on 8 July 2024).