
Scalable and interpretable deforestation detection in the Amazon rainforest

Rodrigo Schuller
IMPA
Rio de Janeiro, Brazil
rodrigo.loro@impa.br

Francisco Ganacim
IMPA
Rio de Janeiro, Brazil
ganacim@impa.br

Paulo Orenstein
IMPA
Rio de Janeiro, Brazil
pauloo@impa.br

Abstract

Deforestation of the Amazon rainforest is a major contributor to climate change, as it is a crucial precipitation regulator, as well as a large natural carbon reserve. While there have been efforts to create real-time algorithms for deforestation detection, they are oftentimes not accurate or interpretable. We leverage multiple input signals, such as satellite imagery, time-series of deforestation indices and scalar measures, to create a single deep learning model that is both interpretable and accurate. We employ a novel dataset with millions of annotated images of the Brazilian Amazon to train our model, as well as class activation mappings to investigate the added value of interpretability in this context.

1 Introduction

Destroying the Amazon rainforest contributes to the global climate crisis by releasing vast amounts of CO₂ into the atmosphere [4]. It also disrupts rainfall patterns, leading to lower agricultural revenues [7]. Understanding the driving forces behind deforestation and precisely quantifying its impacts is important for designing effective public policies. Furthermore, advances in data acquisition and processing, as well as deforestation algorithms, have enabled several initiatives to track deforestation in the Amazon. For example, Instituto do Meio Ambiente e do Homem (Imazon) is an NGO that uses both algorithms and human experts to track deforestation in the Amazon, while the Global Forest Watch project [5] employs real-time algorithm to study and visualize deforestation around the world in near real-time. Many governments also provide their own estimates of deforestation rates, such as Brazil's PRODES [13] and DETER [3].

Imazon's Deforestation Alert System is a particularly accurate source of deforestation data in the Amazon rainforest. This system utilizes change detection algorithms on top of satellite imagery, as well as human supervision, to detect deforestation. It then generates monthly reports analyzing areas affected by illegal fires and other forms of forest degradation. These reports are publicly accessible and play a vital role in combating illegal deforestation and highlighting current trends, which can then be acted upon by local governments.

Their pipeline can be briefly described as follows. Each month, the Normalized Difference Fraction Index (NDFI) indicator [11] is calculated pixelwise across the entire Amazon rainforest using Copernicus Sentinel-2 multispectral imagery. The results are then compared with those from the previous month, and areas showing significant divergence are flagged for further inspection. A team of expert analysts manually reviews these areas, determining whether they represent deforestation events. To make these determinations, various data sources are utilized, including Copernicus Sentinel-2, Landsat 8 and 9 imagery, and occasionally Copernicus Sentinel-1 radar.

In spite of this, tracking deforestation is a challenging and error-prone task. The primary challenge lies in the extensive cloud cover — which is estimated to obscure at least 70% of the Amazon rainforest at any given time — that interfere with satellite readings. Additionally, seasonal events like

floods and receding vegetation over rocky outcrops can create ambiguous situations where vegetation appears to be missing compared to the previous month, despite no deforestation having taken place. The difficulty is compounded by the low resolution of satellite images (10 meters per pixel at best) and the relatively small size of the monthly deforested areas, most of which cover only a few hundred square meters. These factors make deforestation detection highly complex and subjective, further highlighting the importance of auditable and interpretable results.

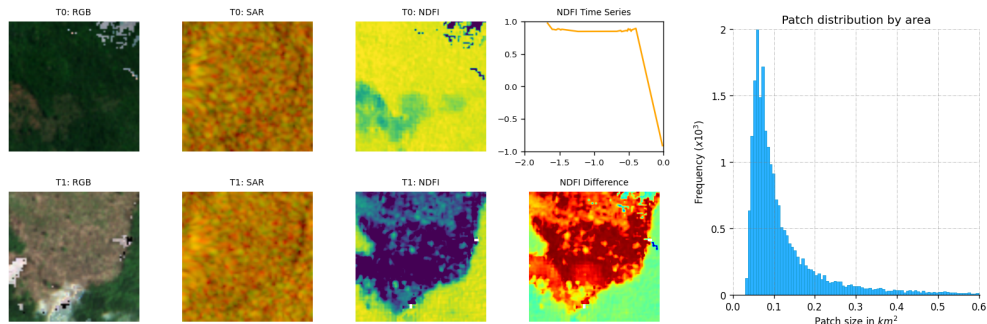
In this work, we describe ongoing efforts to automate the process outlined above through interpretable machine learning algorithms. We utilize a dataset containing hundreds of thousands of manually annotated samples to train several classification networks aimed at detecting deforestation in near real-time. Beyond simply classifying samples, we strive to create a tool that provides insights into the reasons behind the model classifications. To achieve this, we combine images, capturing geographical features, with time series data, reflecting seasonal events, in a compound model. The result is a model with great accuracy and interpretable results. In the following sections, we detail the organization of our dataset and the data processing methods employed; discuss the models used and their training procedures; subsequently, present our results and offer an interpretation of the models' behavior; and, finally, we conclude with our findings and outline next steps in this research.

Related work. Several works focus on deforestation detection, or more generally computing land cover change detection with satellite signals. They range from projects with worldwide scope, such as the full featured monitoring tool Global Forest Watch [5], large datasets as [12]; to specialized studies that focus on high quality, expert labeled datasets, of a single or few biome(s), as in ForestNet [6], that managed to successfully classify driving agents of deforestation on Indonesian tropical forests. The adopted techniques also vary significantly, and include per-pixel applications of pseudo-inverse matrices, known in the remote sensing literature as spectral unmixing, such as the NDFI (2005) [11] and the LCCD [1] indexes; traditional ML algorithms, such as random forests in [10]; and deep learning models, as in the aforementioned ForestNet architecture. To make our models more interpretable, we employ GradCAM [9], an algorithm that generates class activation maps for CNNs.

2 Dataset

Our dataset comprises a collection of polygons outlining patches of land in the Amazon rainforest with high NDFI variance, dating from September 2021 to December 2023. These patches are thus the result of a change detection algorithm that quickly filters for potential deforestation events. Since over 90% of the resulting patches are false positives (i.e., they do not contain actual deforestation), each polygon is then manually reviewed by an expert, who make a final determination as to whether deforestation happened or not. Since the algorithm is heavily skewed towards false positives, we created a smaller, balanced dataset, with 216,000 patches. These were split into train (80%), validation (10%), and test (10%) data folds. Most patches are relatively small, with a mean area of 0.05km^2 . The histogram of the area distribution is shown in Figure 1b.

For each polygon, we gathered the following features: mosaics computed from multispectral satellite images, C-band SAR scans, site statistics (e.g., distance to nearest water basin or river), and the mean NDFI time series over the previous two years. Each single mosaic [2] captures the state of the area patch in a given month, and it is formed by merging all the images available in the month to mitigate issues caused by cloud cover, with each mosaic containing the most recent cloud-free pixels in the sequence. We computed two multispectral mosaics from each data source: one for the period before the potential deforestation and one for the period after. The image sources used were the Copernicus Sentinel-2, Landsat 7, and Landsat 8 satellites. C-band SAR scans, collected from Copernicus Sentinel-1 data, do not suffer from cloud interference, so we used the first scan from the initial period and the last from the subsequent period without needing to compute a mosaic. Site statistics included the terrain elevation model (DEM) from [14] and water distribution information from the JRC Global Surface Water Mapping Layers [8]. Additionally, we computed the mean NDFI time series for each patch over the preceding two years using Sentinel-2 data. Figure 1a shows one sample from our dataset. All data was preprocessed and downloaded via Google Earth Engine.



(a) A sample, with before and after images downloaded from multiple satellites and a time series of the NDFI index, an indicator of vegetation. (b) Patch distribution by area.

Figure 1: Our dataset.

Model	Data Source	Accuracy	Compound Model	Accuracy
ResNet	Sentinel-1	0.865	Baseline ResNet	0.956
ResNet	Sentinel-2	0.958	Baseline ResNet + LSTM	0.957
ResNet	Landsat 7	0.838	Concatenated ResNet	0.954
ResNet	Landsat 8	0.880	Concatenated ResNet + LSTM	0.957
ResNet	DEM	0.663	Logistic Regression	0.962
ResNet	JRC	0.569		
ResNet	Dist to Rivers	0.567		
LSTM	NDFI Time series	0.949		

Table 1: Summary of models tested.

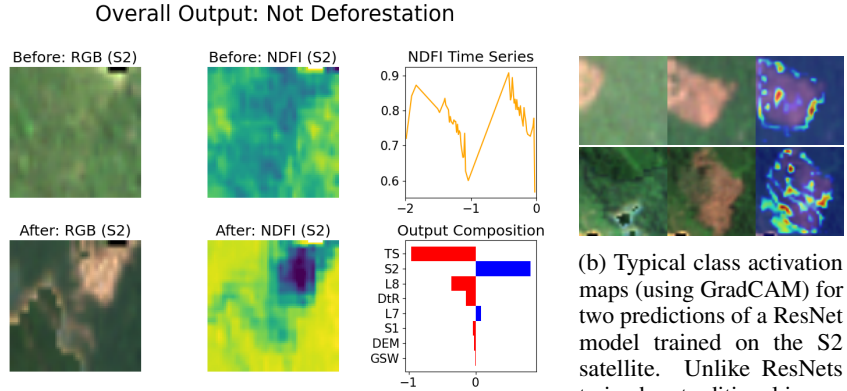
3 Models

To determine if deforestation can be reliably detected using the aforementioned data, we tested a suite of deep learning models. We first evaluated models individually for each data source. For the image data, the best-performing model was a residual CNN followed by linear layers for regression. These models were fed composite data from before and after a candidate deforestation event, stacked as image layers. For the NDFI time series, we considered two architectures: a 1D convolutional model and an LSTM-based model.

We then created a compound model to combine all of the features available. First, we established a baseline by fitting a logistic regression model on the outputs of all the individual models. Next, we retrained this composite network, allowing the gradient to flow back to the individual models. Finally, we created a model by concatenating the flattened outputs of the convolutional features from all the residual networks with the output of the time series model. This combined output was fed to a series of linear layers with activation, followed by a final regression layer. The results are shown in Table 1.

4 Results

By training models with distinct input signals — namely the NDFI time series (TS), and the images Sentinel-2 and 1 (S2/S1), Landsat-7 and 8 (L7/L8), Merit Digital Elevation Model (DEM) Distance to Rivers (DtR), and Global Surface Water (GSW), mapping the location and temporal distribution of surface water — and fitting a logistic regression on their outputs, we are able to quantify the contributions from each data type. Figure 2a showcases one of such predictions, in which the negative signal from the time series of the NDFI index overcomes the positive contribution from the before and after Sentinel-2 images, yielding a correct negative prediction. This pattern suggests a seasonal degradation, which can be confirmed by the NDFI Time Series plot. The current accuracy of 96% (see Table 1) is enough to automate several steps on Amazon’s pipeline, and to develop a near real-time monitoring system.



(a) Contributions per input signal by the logistic regression. Note the positive contribution from the before and after images, and the negative one from the NDFI time series, an index that aims to measure vegetation. This composition pattern suggests a seasonal degradation event, which we can confirm by inspecting the NDFI Time Series plot.

(b) Typical class activation maps (using GradCAM) for two predictions of a ResNet model trained on the S2 satellite. Unlike ResNets trained on traditional image datasets, the activated areas tend to concentrate on the edges of the deforested areas.

Figure 2: Interpretability due to logistic regression (left) and GradCAM (right).

Besides unravelling the predictions with respect to the different input signals, we used GradCAM on the CNN models to compute the importance of features with respect to the (x, y) coordinates. Typically, the activations concentrate on the edges of the deforested area, as shown in Figure 2b, which is different from the usual GradCAM results when applied to ResNet on RGB images. These results point to future research direction: either mask-based data augmentation or changes in the last pooling layer can fill up the deforested areas in the class activation heatmaps, or the attention on the borders could be an indicator of an intrinsic characteristic of the dataset.

5 Conclusion and Future Work

Although GradCAM and logistic regression show promising interpretability insights, specialized multimodal architectures could further improve our results, both with respect to predictive performance and accurate unsupervised classification of degradation events (seasonality, degradation agent, morphological structure, etc).

Particular characteristics of the main family of satellites utilized in our study, Sentinel-2, are a fertile ground for interpretability inquiries. One of them is the temporal delay between recordings of different bands, which causes parallax artifacts on clouds and other atmospheric phenomena, creating frequent rainbow-like structures, which in turn tends to reduce the effectiveness of per-pixel land cover change detection algorithms. Understanding how the CNN models can achieve high accuracy on a biome full of clouds such as the Amazon rainforest, despite of these artifacts, is a natural step towards the development of a specialized architecture. Finally, we would also like to explore applications of conformal prediction in our problem, both as a way to provide uncertainty quantification and to benchmark our results against state-of-the-art alternatives, in terms of large-scale temporal prediction accuracy stability.

References

- [1] Eric L. Bullock, Curtis E. Woodcock, and Pontus Olofsson. Monitoring tropical forest degradation using spectral unmixing and Landsat time series analysis. *Remote Sensing of Environment*, 238:110968, March 2020.
- [2] Jeffrey A. Cardille, Morgan A. Crowley, David Saah, and Nicholas E. Clinton, editors. *Cloud-Based Remote Sensing with Google Earth Engine: Fundamentals and Applications*. Springer International Publishing, Cham, 2024.

- [3] Cesar Guerreiro Diniz, Arleson Antonio de Almeida Souza, Diogo Corrêa Santos, Mirian Correa Dias, Nelton Cavalcante da Luz, Douglas Rafael Vidal de Moraes, Janaina Sant'Ana Maia, Alessandra Rodrigues Gomes, Igor da Silva Narvaes, Dalton M. Valeriano, Luis Eduardo Pinheiro Maurano, and Marcos Adami. DETER-B: The New Amazon Near Real-Time Deforestation Detection System. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 8(7):3619–3628, July 2015. Conference Name: IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing.
- [4] Luciana V. Gatti, Luana S. Basso, John B. Miller, Manuel Gloor, Lucas Gatti Domingues, Henrique L. G. Cassol, Graciela Tejada, Luiz E. O. C. Aragão, Carlos Nobre, Wouter Peters, Luciano Marani, Egidio Arai, Alber H. Sanches, Sergio M. Corrêa, Liana Anderson, Celso Von Randow, Caio S. C. Correia, Stephane P. Crispim, and Raiane A. L. Neves. Amazonia as a carbon source linked to deforestation and climate change. *Nature*, 595(7867):388–393, July 2021.
- [5] M. C. Hansen, P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov, L. Chini, C. O. Justice, and J. R. G. Townshend. High-Resolution Global Maps of 21st-Century Forest Cover Change. *Science*, 342(6160):850–853, November 2013.
- [6] Jeremy Irvin, Hao Sheng, Neel Ramachandran, Sonja Johnson-Yu, Sharon Zhou, Kyle Story, Rose Rustowicz, Cooper Elsworth, Kemen Austin, and Andrew Y. Ng. ForestNet: Classifying Drivers of Deforestation in Indonesia using Deep Learning on Satellite Imagery, November 2020. arXiv:2011.05479 [cs, eess].
- [7] Argemiro Teixeira Leite-Filho, Britaldo Silveira Soares-Filho, Juliana Leroy Davis, Gabriel Medeiros Abrahão, and Jan Börner. Deforestation reduces rainfall and agricultural revenues in the Brazilian Amazon. *Nature Communications*, 12(1):2591, May 2021. Publisher: Nature Publishing Group.
- [8] Jean-François Pekel, Andrew Cottam, Noel Gorelick, and Alan S. Belward. High-resolution mapping of global surface water and its long-term changes. *Nature*, 540(7633):418–422, December 2016.
- [9] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. *International Journal of Computer Vision*, 128(2):336–359, February 2020. arXiv:1610.02391 [cs].
- [10] Carlos M. Souza, Luis A. Oliveira, Jailson S. De Souza Filho, Bruno G. Ferreira, Antônio V. Fonseca, and João V. Siqueira. Landsat sub-pixel land cover dynamics in the Brazilian Amazon. *Frontiers in Forests and Global Change*, 6:1294552, December 2023.
- [11] Carlos M. Souza, Dar A. Roberts, and Mark A. Cochrane. Combining spectral and spatial information to map canopy damage from selective logging and forest fires. *Remote Sensing of Environment*, 98(2-3):329–343, October 2005.
- [12] Radost Stanimirova, Katelyn Tarrío, Konrad Turlej, Kristina McAvoy, Sophia Stonebrook, Kai-Ting Hu, Paulo Arévalo, Eric L. Bullock, Yingtong Zhang, Curtis E. Woodcock, Pontus Olofsson, Zhe Zhu, Christopher P. Barber, Carlos M. Souza, Shijuan Chen, Jonathan A. Wang, Foster Mensah, Marco Calderón-Loor, Michalis Hadjidakou, Brett A. Bryan, Jordan Graesser, Dereje L. Beyene, Brian Mutasha, Sylvester Siame, Abel Siampale, and Mark A. Friedl. A global land cover training dataset from 1984 to 2020. *Scientific Data*, 10(1):879, December 2023.
- [13] Dalton M Valeriano, Eliana MK Mello, José Carlos Moreira, Yosio E Shimabukuro, Valdete Duarte, IM Souza, JR Santos, Claudio CF Barbosa, and RCM Souza. Monitoring tropical forest from space: the PRODES digital project. *International Archives of Photogrammetry Remote Sensing and Spatial Information Sciences*, 35:272–274, 2004.
- [14] Dai Yamazaki, Daiki Ikeshima, Ryunosuke Tawatari, Tomohiro Yamaguchi, Fiachra O’Loughlin, Jeffery C. Neal, Christopher C. Sampson, Shinjiro Kanae, and Paul D. Bates. A high-accuracy map of global terrain elevations. *Geophysical Research Letters*, 44(11):5844–5853, June 2017.