Near Real-Time Detection of EVI Time-Series Breakpoints Using Bayesian Inference for Deforestation Monitoring in the Chaco Forest

Francisco Grings[1,2], Francisco González Bianco[2,3], Esteban Roitberg[2,4,5], Natalia Morandeira[2,6], Javier Arellana[1,2,8], Maira Gayol[2,6,7]

[1] Instituto de Astronomía y Física del Espacio (CONICET-UBA), CABA, Argentina
[2] Pixel - Satellite-based Environmental Data Analysis
[3] Centro de Estudios sobre Patrimonios y Ambiente, Universidad Nacional de San Martín (UNSAM), San Martín, Buenos Aires, Argentina
[4] Instituto Tecnológico de Buenos Aires, CABA, Argentina
[5] Escuela de Ciencia y Tecnología, Universidad Nacional de San Martín, Buenos Aires, Argentina
[6] Instituto de Investigación e Ingeniería Ambiental, CONICET-UNSAM, Escuela de Hábitat y Sostenibilidad, Universidad Nacional de San Martín, General San Martín, Buenos Aires, Argentina
[7] Comisión de Investigaciones Científicas de la Provincia de Buenos Aires, Argentina
[8] Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física. Buenos Aires, Argentina

Keywords: Deforestation, Bayesian inference, Chaco Forest, real-time.

Abstract

Deforestation poses a significant threat to natural ecosystems, particularly in Argentina's Chaco region—one of the world's most rapidly changing forest areas. This study focuses on the detection of sudden deforestation events, where forest cover is rapidly removed within a few months. Monitoring such changes across vast areas requires the use of satellite-based vegetation indices, such as the Enhanced Vegetation Index (EVI) and Normalized Difference Vegetation Index (NDVI) from MODIS. However, accurately identifying deforestation events is challenging due to seasonal variability, sensor noise, data gaps, and algorithmic inconsistencies. These factors can obscure true deforestation signals or generate false positives. To address these issues, a robust detection approach must explicitly model time-series dynamics, capturing trends, seasonality, and uncertainty, to reliably distinguish genuine deforestation breakpoints from natural variation and noise. In this paper, three models for the detection of breakpoints in EVI time series were proposed: a simple z-score anomaly detector, and two fully Bayesian models; one temporally uncorrelated and one fully correlated. Results indicate that the Bayesian schemes significantly improve over the naive approach (zscore: AUC=0.921, F1-score=0.870, Bayes: AUC=0.959, F1-score=0.925), for a reasonable cost in computing time ×1000.

1. Introduction

Deforestation is one of the biggest dangers to natural forests today, in particular for the Chaco region of Argentina, which is one of the fastest-changing forest areas in world. Most of the land being cleared is used to grow crops like soybeans or to raise cattle, because of the high demand from around the world.

To keep track of how forests are being cleared in a huge area like the Argentine Chaco, the use of satellite images and time series analysis are mandatory (Grings et al. (2020)). This is important because deforestation doesn't always happen the same way—it can be fast or slow. There are two main types of deforestation based on how quickly it happens:

1. **Sudden Deforestation** – This happens when a large area of forest is quickly cleared, (sometimes but not always replaced with pastures or crops).

2. **Gradual Deforestation (Forest Degradation)** – This takes longer and causes the forest to become patchy or damaged before it's eventually removed.

In this study, we're focusing only on the fast kind—when the forest is totally removed in just a few months. This kind of deforestation can happen any time of the year, but it's especially hard to detect during the winter. That's because the satellite images used show smaller changes in vegetation during colder months, because seasonal vegetation dynamics reduce canopy greenness, making it tougher to spot the damage.

Satellite-derived time series of vegetation indices such as EVI and NDVI from MODIS (Moderate-Resolution Imaging Spectroradiometer) are widely used for deforestation monitoring. Although MODIS has a coarse spatial resolution (250 m), limiting the detection of small clearings, it offers unique advantages: more than two decades of continuous coverage, high temporal resolution, and well-validated vegetation indices. These properties make MODIS particularly suitable for testing Bayesian breakpoint detection methods in long time series. In contrast, the Harmonized Landsat-Sentinel (HLS) dataset provides finer spatial resolution (30 m) and increased acquisition frequency but has only been available since 2013, making it unsuitable for long historical analyses. However, HLS represents a promising source for future applications focused on recent years. It is also important to consider the typical size of deforested areas. Vallejos et al. (2015) report an average patch size of 61.99 $\pm$ 0.29 hectares, indicating that the MODIS pixel size is sufficiently granular to capture most events. While HLS could produce more detailed alerts, it also entails significantly higher computational demands. Finally, the daily acquisition frequency of MODIS minimizes the impact of cloud cover on the data record.

Deforestation typically appears as a sharp drop or "breakpoint" in these index values, as dense forest canopy is removed and

https://doi.org/10.5194/isprs-annals-X-3-W4-2025-191-2026

only bare ground or sparse vegetation remains. However, reliably detecting such breakpoints is challenging due to several factors:

- Seasonal and Phenological Variability: Vegetation indices fluctuate seasonally, which can mask or mimic deforestation signals. Distinguishing a genuine forest loss event from normal seasonal drops requires accounting for seasonal patterns.
- Data Gaps and Noise: Noise from atmospheric effects or sensor issues can introduce data gaps. For example, the MODIS VI product accuracy is about ±0.015 for EVI under ideal conditions, but errors can increase to 0.04–0.10 in index units when data quality is poor (MODIS Land Science Team (2024)).
- Algorithm Inconsistencies: The MODIS EVI algorithm itself can switch modes under certain conditions, introducing artificial jumps. Notably, when a pixel is very bright (e.g. clouds or snow), the standard 3-band EVI is replaced with a 2-band EVI (excluding the blue band) to avoid overcorrection (Center for Tropical Agriculture and Human Resources (CTAHR) (2024)). Earlier product versions even fell back to a Soil-Adjusted VI in such cases (Center for Tropical Agriculture and Human Resources (CTAHR) (2024)). These algorithm changes between acquisition dates can cause small discontinuities in the time series that are unrelated to actual vegetation change. A robust detection method should not mistake these artifacts for deforestation.

Given these challenges, a detection approach must explicitly model the time-series structure (trend, seasonality, noise) and account for uncertainties, so that genuine deforestation breakpoints can be distinguished from normal variation and noise.

2. Background

2.1 Basics

A number of time-series change detection techniques have been applied to remote sensing data in the past. One well-known approach is Breaks For Additive Season and Trend (BFAST) (Verbesselt et al. (2010)), which decomposes a time series into trend, seasonal, and remainder components and iteratively finds breakpoints in those components. Other widely used methods include LandTrendr (segmented regression of Landsat time series (Kennedy et al. (2010))) and the CCDC algorithm (Continuous Change Detection and Classification (Zhu and Woodcock (2014))), which have been used to map forest loss by fitting temporal segments to long-term satellite data. These classical methods, however, are generally frequentist in nature – they provide point estimates of change timing and often require ad-hoc thresholds or significance tests to decide if a change is real. They may not fully quantify the uncertainty in the detected breakpoints, and they can be sensitive to outliers or noise if not carefully tuned.

2.2 Bayesian Approaches for Breakpoint Detection

A Bayesian approach treats the time series within a probabilistic framework. In a Bayesian breakpoint model, each observation is treated as a random variable generated from some underlying state (e.g. "forest" or "deforested" regime), and one assigns prior probabilities to the existence and timing of breakpoints. The result of a Bayesian analysis is not just a single change date, but a posterior probability distribution over possible breakpoints – effectively measuring how likely a deforestation event is at each time, given the data and prior expectations (Barraza and Grings (2016)). This framework naturally accounts for uncertainties and allows incorporating temporal correlation and prior knowledge (for example, an expectation that deforestation is infrequent at a given date or area).

An example of a recent application of this framework is Wendelberger et al. (2021), which extends the standard Bayesian Online Changepoint Detection (BOCPD) algorithm of Adams (2007) by incorporating a seasonal regression model and an outlier-tolerant likelihood (Wendelberger et al. (2021)). In their approach, each new observation is modeled via a multivariate linear regression (capturing seasonal cycles and possibly multiple spectral indices), and the algorithm continuously updates the probability of a breakpoint.

A related development is the Bayesian ensemble approach for time-series decomposition, exemplified by the BEAST algorithm by Zhao et al. (2019) (Zhao et al. (2019)). BEAST is a general Bayesian time series model that simultaneously estimates trend, seasonal components, and breakpoints by averaging over many possible models (e.g. like an ensemble BFAST). Unlike deterministic methods which pick a single best decomposition, BEAST uses Bayesian model averaging to combine many competing models of different complexity. This means it can quantify the uncertainty in the number and timing of breakpoints and in the shape of trend/seasonal components.

The authors also published several approaches to estimate breakpoints in satellite time series to monitor deforestation events in Chaco Forest. These works are relevant approaches to contextualize this work. The first was based on an XGBoost approach (Barraza et al. (2020)). In the second, we included Land Surface Temperature (LST) data to evaluate its relevance in breakpoint detection (Roitberg et al. (2018)). Finally, we also played with the idea of using convolutional neural networks to model the time series, obtaining promising results published in (Grings et al. (2020)).

3. Methodology

In this study, we propose three approaches to implement a breakpoint detection algorithm to detect deforestation from EVI time series data in Chaco Forest:

1. A basic z-score ($z\text{-}score = (x - mean(x))/var(x)$) approach, in which we inform an alert when the z-score for a given date is below a given threshold (that can be tuned to maximize a given metric). Since the z-score automatically takes into account date mean values and variances, we selected this model as a benchmark case.

2. An approach based on a free-form, fully parametric scheme that models each time series date as a random variable, first introduced in (Gregory and Loredo (1992)). This approach can be formulated using uncorrelated random variables for every date (Bayes-naive) and fully correlated dates (Bayes-full).

3.1 A periodic signal of unknown shape and period: mathematical formulation

The approach developed by (Gregory and Loredo (1992)) was to assume that our unknown periodic time series can be modeled as a succession like

$$X = (x_0, x_1, .., x_N, x_0, x_1)$$

in which x_i is a random variable -that corresponds to a given date- and is in principle correlated with all the other x_i. An example of this approach is presented in Fig. 1.

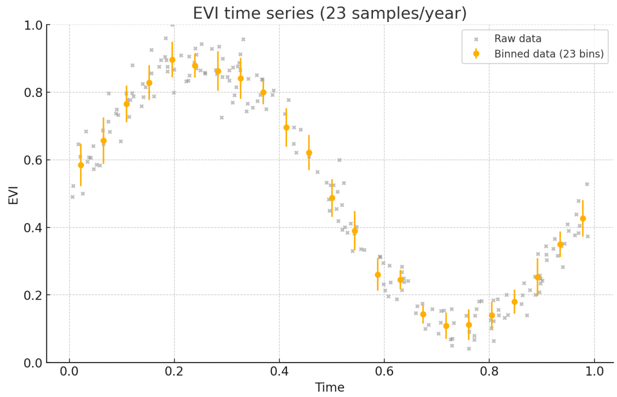

Figure 1. A scheme of the Bayesian model presented in (Gregory and Loredo (1992)) for periodic time series of unknown shape.

Once given this construction scheme, we only need to provide prior distributions for the x_i and a likelihood function. For the simpler case (Bayes-naive), we assume the priors to be normally distributed with no prior correlation between them [1]. In the more advanced approach (Bayes-full), we also assumed a normal distribution for every x_i, but also explicitly modeled the prior correlation among them with a full covariance matrix C of size $N \times N$. This later approach is in principle capable of not only learn the mean values of the dates, but also of using these learned mean values of other dates to help the inference in any other given date.

Finally, a normal likelihood was adopted. This is based on the assumption that the MODIS EVI product is calibrated, implying that the mean residual $(\mathrm{EVI} - \mathrm{model})$ is zero. The noise was characterized by a normal distribution with variance $\sigma_{\mathrm{EVI}} = 0.025$ [2].

Once the model is defined, model parameters need to be estimated from real data. To this end, we defined a period of time for which the forest was undisturbed (calibration period). In this calibration period, the posterior distribution for all model parameters was obtained using PyMC (PyMC Developers (2024b)). In the PyMC framework, after the model is calibrated with the available data, the posterior value of the next date $EVI_{(t+1)}$ is computed using the posterior predictive, in which both the posterior of the parameters and the observation error are taken into account (PyMC Developers (2024a)). These values are compared with the observed values $(EVI(obs)_{(t+1)})$ and its probability in the context of the model is then computed.

3.2 Deforestation “Alert” criteria

We defined an “Alert” for each model according to its algorithm. However, during the evaluation, we observed inconsistencies with isolated alerts on the initial evaluation date. To address this, we implemented a ”double alert” criterion: an alert is induced only when two consecutive evaluations produce an alert. This approach reduced false positives that stemmed from such inconsistencies.

[1] This only implies no prior correlation is assumed. The posteriors could be strongly correlated

[2] as discussed in section 1, the real structure of the noise is very complex, and depends on the specific area and date of acquisition

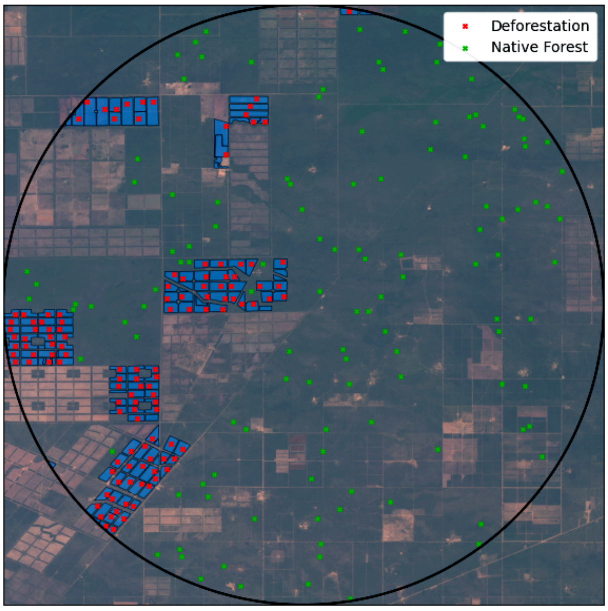

Figure 2. Study area located at the border between Chaco and Salta provinces, Argentina. Light blue polygons indicate plots deforested during the period considered; red dots represent MODIS pixels with known deforestation dates inside those plots, and green dots correspond to undisturbed native forest.

4. Study area

The study area is located within the Semiarid Chaco subregion of the Argentine Chaco Forest, covering approximately 250,000 hectares. This region is characterized by mixed xerophilous forests, predominantly composed of deciduous or semideciduous species adapted to moderate rainfall stress, with annual precipitation ranging between 500 and 800 mm. Dominant species include Aspidosperma quebracho-blanco (white quebracho) and Schinopsis lorentzii (red quebracho), which form part of the subtropical seasonal forest landscape.

Previous research in this and related areas resulted in a GIS database identifying the location and monthly timing of deforestation events in various native forest regions. A second dataset used in this study consists of remote sensing time series derived from the MODIS MOD13Q1 product, spanning from February 2000 to February 2020, with a temporal resolution of 16 days and a spatial resolution of 250 meters. These time series are categorized into two classes: Deforested (2,230 instances), each associated with a precise monthly deforestation date between 2002 and 2016 based on the GIS data described above; Native Forest (22,404 instances), representing undisturbed native forest.

Each time series includes vegetation indices (NDVI and EVI), four spectral bands (red, near-infrared, blue, and mid-infrared), and associated quality indicators. The dataset was developed to support near-real-time deforestation detection and broader environmental monitoring. A detailed description of these databases can be found in (Roitberg et al. (2021)).

https://doi.org/10.5194/isprs-annals-X-3-W4-2025-191-2026 | © Author(s) 2026. CC BY 4.0 License.

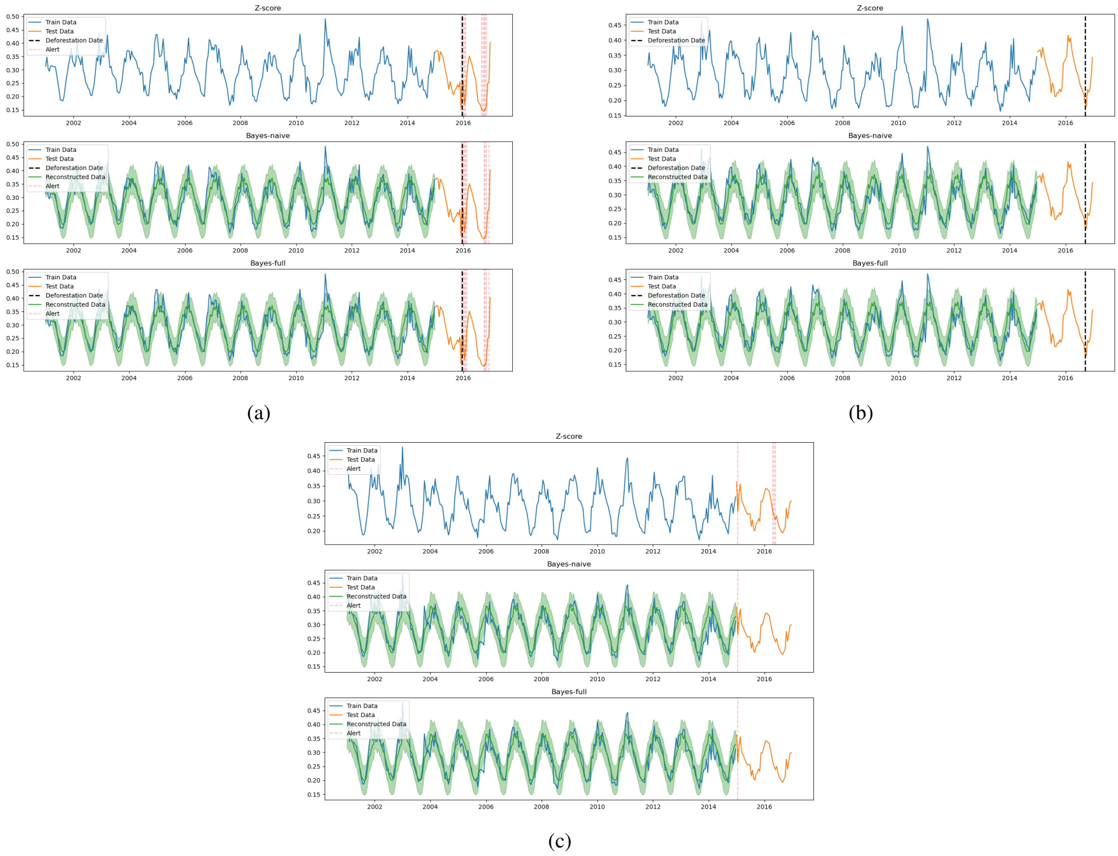

Figure 3. Example of three EVI time series for a given deforested pixel and correctly detected (true positive, top left), a deforested pixel but not detected (false negative, top right), and native forest pixel incorrectly detected as deforested (false positive, bottom). In each panel, first row is z-score model, the second Bayes-naive and the third Bayes-full. Training period is indicated in green. Detected breakpoints are marked as red vertical lines.

For this manuscript, only polygons deforested during 2015 and 2016 were included. For each polygon, a time series was constructed for every pixel; however, due to the high similarity among them, we selected a single representative pixel time series per polygon to avoid redundancy. This procedure resulted in 122 deforested time series. An equal number of time series from the Native Forest class was randomly sampled for comparison.

5. Results

The three algorithms were applied to all EVI time series in the study area, including pixels with observed deforestation and pixels of undisturbed forest (both conditions identified using high-resolution imagery). During the calibration period, each model estimated the posterior probability distribution of its parameters. Then, as described in Section 3.1, a posterior predictive distribution was computed to estimate the probability density function (pdf) of $EVI_{(t+1)}$, and subsequently the probability of the observed EVI value within the model context. Examples are shown in Fig. 3, which presents deforested and undisturbed time series to illustrate model behavior.

As shown, all three models can detect deforestation events, labeled as “Alerts,” corresponding to breakpoints in the time series (Fig. 3, top left). This case illustrates a true positive.

Fig. 3, top right, shows a false negative: a deforested pixel not detected by any of the models. Although the observed EVI decreases at the deforestation date, the magnitude of the drop falls within the range of previous fluctuations. Finally, Fig. 3, bottom, shows a false positive, where alerts are generated over an undisturbed forest series. Notably, all three models trigger an alert on the same date, but only the Z-score model produces a sequence of successive alerts at posterior dates, underscoring the stronger performance of the Bayesian approaches.

An comparison of performance of the three models is made by computing the Receiver Operating Characteristic (ROC) for the three models in Fig. 4. It can be seen that, though all models performed well in this task, the Bayesian schemes (Bayes-naive and Bayes-full) present marginally better AUC results. This is particularly important for low threshold values, for which high values of True Positive Rate (TPR) can be obtained for relatively low values of False Positive Rate (FPR), which is very important in an operational scheme.

Finally, a summary of the results for all algorithm and pixels is presented in Table 1. As seen, both Bayesian models presented significantly better F1-scores and accuracies. However, the more complex model Bayes-full is much more computationally expensive, hindering its usability in real-time applications.

6. Discussion

In this paper, three models for the detection of breakpoints in EVI time series were proposed. The models can be seen as an increasingly complex approximation to a similar approach. In

Model	Threshold	TP	FN	FP	TN	Accuracy	Precision	Recall	FPR	F1 Score	AUC-ROC	lag_05q	lag_50q	lag_95q	Time
Z-score	0.008	111	11	22	100	0.865	0.835	0.910	0.180	0.871	0.922	-62.0	8.0	426.0	1
Bayes-naive	0.001	118	4	15	107	0.922	0.887	0.967	0.123	0.925	0.959	-119.0	20.5	159.7	1000
Bayes-full	0.001	116	6	15	107	0.914	0.885	0.951	0.123	0.917	0.957	-119.0	20.5	144.5	100000

Table 1. Comparison of different models across multiple performance metrics

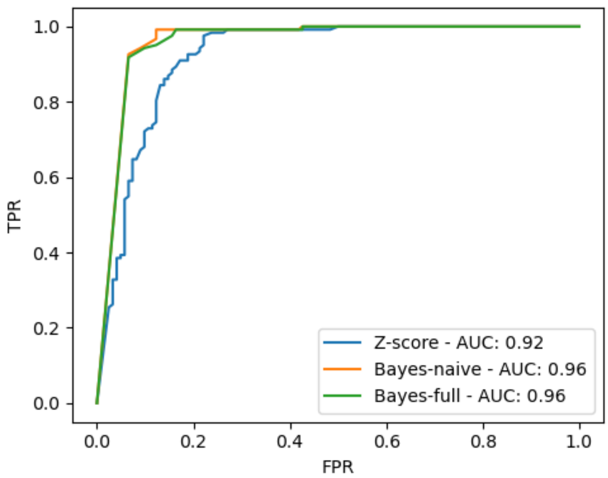

Figure 4. ROC AUC (Area Under the Curve) curves for the three models

the first model, an anomaly is defined as a strong enough deviation of the normalized time series (z-score). In the second one, date's distributions are explicitly modeled, in order to obtain posterior distribution for every date to compare it with future unobserved values, and an anomaly is defined as a p-value over a posterior distribution (Bayes-naive). In the last one, a single N-dimensional distribution is proposed for the entire periodic time series, with N mean values and a correlation matrix of $N \times N$.

The more simple model presented an acceptable performance given its simplicity, proving that it is a robust approach to this problem. The first Bayesian model presents better metrics, mainly due to its ability to include a distribution to every date instead of just mean values and variances. However, this increase comes at a cost of $\times 1000$ overhead in computing time. Finally, the most complex Bayesian approach fails to improve significantly over the simpler model. This result is a little surprising, but is probably related to the fact that in its original application (Gregory and Loredo (1992)), much larger time series were used for calibration, which included several periods of the unknown signal. Although a long period for a remote sensing time series, we just used 10 years of data for calibration.

As a future approach, it is possible to better model data inter-annual variability, which in the analyzed models is only "seen" as over-dispersed posterior distribution for a given date. However, care must be taken in order to use harmonic+trend+noise models like BFast, since they are very sensitive to noise and produce a lot of false positives. A more robust approach needs to be devised that robustly takes into account natural inter-annual variability.

References

Barraza, V., Grings, F., 2016. Detection of trend change-point in passive microwave and optical time series using bayesian inference over the dry chaco forest. *Proceedings*, 1number 2, MDPI, 45.

Barraza, V., Roitberg, E., Grings, F., 2020. Monthly deforestation prediction using extreme gradient boosting based on modis evi time series. *2020 IEEE Congreso Bienal de Argentina (ARGENCON)*, IEEE, 1–5.

Center for Tropical Agriculture and Human Resources (CTAHR), 2024. Mod13: Vegetation index user's guide – section 1. introduction. `https://www.ctahr.hawaii.edu/grem/mod13ug/sect0001.html`. Accessed: 2025-06-09.

Gregory, P., Loredo, T. J., 1992. A new method for the detection of a periodic signal of unknown shape and period. *Astrophysical Journal, Part 1 (ISSN 0004-637X), vol. 398, no. 1, Oct. 10, 1992, p. 146-168., Research supported by NSERC*, 398, 146–168.

Grings, F., Roitberg, E., Barraza, V., 2020. EVI Time-Series Breakpoint Detection Using Convolutional Networks for Online Deforestation Monitoring in Chaco Forest. *IEEE Transactions on Geoscience and Remote Sensing*, 58(2), 1303–1312. http://dx.doi.org/10.1109/TGRS.2019.2945719.

Kennedy, R. E., Yang, Z., Cohen, W. B., 2010. Detecting trends in forest disturbance and recovery using yearly Landsat time series: 1. LandTrendr—Temporal segmentation algorithms. *Remote Sensing of Environment*, 114(12), 2897–2910.

MODIS Land Science Team, 2024. Modis vegetation indices (mod13) validation status. `https://modis-land.gsfc.nasa.gov/ValStatus.php?ProductID=MOD13`. Accessed: 2025-06-09.

PyMC Developers, 2024a. Posterior predictive sampling — pymc documentation. `https://www.pymc.io/projects/docs/en/stable/learn/core_notebooks/posterior_predictive.html`. Accessed: 2025-06-09.

PyMC Developers, 2024b. Pymc. `https://github.com/pymc-devs/pymc`.

Roitberg, E., Barraza, V., Grings, F., Salvia, M., Perna, P., Barber, M., 2018. Near real time multisensor algorithm for deforestation alert over the dry chaco forest. *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 818–821.

Roitberg, E., Barraza, V., Salvia, M., Grings, F., 2021. Remote sensing time series database for deforestation detection. *VII Simposio Argentino de Ciencia de Datos y GRANdes DAtos (AGRANDA 2021)-JAIIO 50 (Modalidad virtual)*.

Vallejos, M., Volante, J. N., Mosciaro, M. J., Vale, L. M., Bustamante, M. L., Paruelo, J. M., 2015. Transformation dynamics of the natural cover in the Dry Chaco ecoregion: a plot level geo-database from 1976 to 2012. *Journal of Arid Environments*, 123, 3–11.

Verbesselt, J., Hyndman, R., Zeileis, A., Culvenor, D., 2010. Phenological change detection while accounting for abrupt and gradual trends in satellite image time series. *Remote Sensing of Environment*, 114(12), 2970–2980.

Wendelberger, L. J., Gray, J. M., Reich, B. J., Wilson, A. G., 2021. Monitoring Deforestation Using Multivariate Bayesian Online Changepoint Detection with Outliers. *arXiv preprint arXiv:2112.12899*.

Zhao, K., Wulder, M. A., Hu, T., Bright, R., Wu, Q., Qin, H., Li, Y., Toman, E., Mallick, B., Zhang, X. et al., 2019. Detecting change-point, trend, and seasonality in satellite time series data to track abrupt changes and nonlinear dynamics: A Bayesian ensemble algorithm. *Remote sensing of Environment*, 232, 111181.

Zhu, Z., Woodcock, C. E., 2014. Continuous change detection and classification of land cover using all available Landsat data. *Remote sensing of Environment*, 144, 152–171.