HFLD Research

Technical Report

Predicted Deforestation Rates Statistical Modeling

Produced for the Singapore National Climate Change Secretariat (NCCS)

22 May, 2026

Incentivizing investment in real climate action

Sylvera

Details

Prepared for

Singapore Prime Minister's Office (PMO), National Climate Change Secretariat (NCCS)
100 High Street, #03-01 The Treasury
Singapore, 179434

Prepared by

Sylvera Ltd
18-29 Mora St, London
United Kingdom, EC1V 8BT

Executive summary

Context

Singapore's National Climate Change Secretariat (NCCS) commissioned Sylvera to undertake statistical modeling work for predicting deforestation rates for HFLD jurisdictions. This technical report presents the results from two statistical models tested as well as the underlying data sources and statistical tests conducted.

Key findings

The number of HFLD jurisdictions has declined over time. The data shows that the number of jurisdictions meeting HFLD status has decreased over time, and those that have lost their status show accelerating deforestation rates. However, jurisdictions that have maintained their status demonstrate persistently low deforestation rates, indicating that a meaningful subset of HFLDs will remain at low risk. This heterogeneity in risk profiles gives credence to the use of classification methods such as the logistic regression employed in Model 2 to differentiate between jurisdictions.

The historical average fails to capture forward-looking risk. The classic JREDD baseline approach — using historical average deforestation rates — is unable to capture observed spikes in deforestation and therefore cannot predict elevated risk for jurisdictions where deforestation has not historically been high. This is a fundamental limitation for the construction of a representative baseline of emissions in a dynamic risk environment.

Model 1 (Single-stage Random Forest model) best captures variability; Model 2 (Two-Step Logistic Regression and Random Forest) offers a more conservative alternative. Model 1 is more accurate at predicting deforestation rates across the 2015–2020 test period, effectively capturing both variance and extreme values as reflected in its lower MAE and RMSE. Model 2 produces more conservative predictions clustered at lower deforestation rates, making it more suitable for jurisdictions with historically low rates where overprediction carries greater consequences.

The Logistic regression from Model 2 and associated risk classification can serve as a standalone screening tool. The transition probabilities produced by the logistic regression step of Model 2 can be used independently as an initial screening tool to identify HFLD jurisdictions at risk of losing their status within a given time window.

Country-specific and spatially-explicit models are needed for improved accuracy. The variation in performance across models indicates that country-specific approaches may be necessary for more accurately predicting deforestation rates at the individual jurisdiction level. The statistical models examined in this report also lack spatial context — specifically where within a jurisdiction deforestation is likely to occur — which is a key consideration for predicting future forest loss. Modelled results perform adequately given the noisy nature of deforestation data but should not be relied upon for crediting purposes without further refinement.

In conclusion, HFLD ecosystems are precious, and there is a **role for carbon markets** to play in ensuring their protection, provided that appropriate assessments of environmental integrity can be made.

An assessment of environmental integrity should be made at the **jurisdiction- or country-specific level**.

Table of contents

Introduction

Background and scope

Singapore's National Climate Change Secretariat (NCCS) commissioned Sylvera to undertake statistical modeling work for predicting deforestation rates for HFLD jurisdictions. This technical report presents the results of the statistical modeling work including the results from the two models tested as well as the underlying data sources and statistical tests conducted.

More specifically, this technical report sets out the results for two sets of statistical models:

- Model 1: A global single-stage random forest model trained on a global set of HFLD's
- Model 2: A two-step HFLD logistic regression and random forest model applied to a global set of HFLD's

The models are trained on a set of covariates—such as socioeconomic, environmental, and deforestation patterns—across many jurisdictions. They learn patterns from this data and then output predicted deforestation rates for HFLD jurisdictions. These predictions can reflect counterfactual risk, or the level of deforestation that would typically be expected based on the 19 covariate values, rather than observed outcomes. Model 1 is trained on 2,570 jurisdiction-year observations, while Model 2 is trained on HFLD jurisdiction-year observations across each year from 2005 to 2014 for the logistic step, and 160 jurisdiction-year observations for the Random Forest step. The difference in Random Forest training was that Model 2 trained only on HFLD's which transitioned, reducing the number of observations. Deforestation rates were derived using Hansen et al.'s Global Forest change data, which provides global estimates of deforestation annually from 2001.

A statistical model looks at many observations where both the inputs (covariates) and the outcome (ex. deforestation) are known, and 'learns' from patterns in the data. Then, when given the covariates for an area, it can predict the outcome based on what it learned.

Figure 1 below represents a single example of input covariates and their associated variable of interest (deforestation). The covariate values in this example are from Bengo, Angola in the year 2002, which relate to a known deforestation rate of 1%. Based on thousands of these inputs, the models learn from patterns in the data and then can produce predictions, based only on input covariates.

Observed data example		
Covariates	**Value**	**Unit**
Mean patch size	0.01	ha
Edge density	0.47	m/ha
Elevation	237.51	m (ASL)
GDP PPP	3,473.17	$USD
HDI	0.41	0-1 scale
Mining %	0.00	% of total area with mining activity
Nightlight	0.01	Mean DN (0–63)
Patch density	0.21	n / 100 ha
Plantation %	0.00	% of total area with forest plantations
Pop.density	10.18	People per km²
Precipitation	631.55	Mean annual (mm)
Road density	0.19	km of road per km²
Rural pop. %	47.54	% of total population living in rural areas
Slope	0.71	Degrees (°)
Temperature	295.76	Annual mean (K)
Forest area %	82.87	% of total area forested
Deforestation	**1%**	

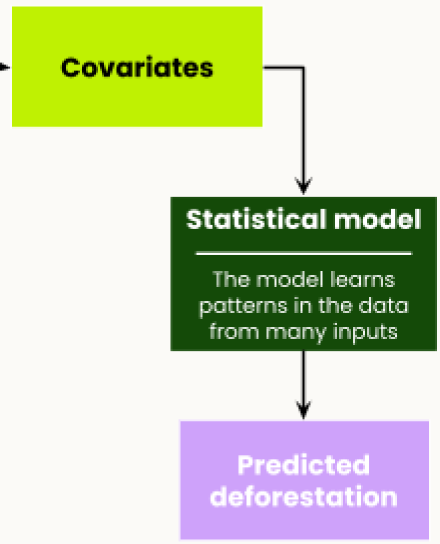

Figure 1. Conceptual diagram of a statistical model
Source: Sylvera, 2026

This report

This technical report presents findings from Sylvera's statistical model for predicting deforestation rates for HFLD jurisdictions. More specifically, it presents two model designs - a single stage random forest model and a two step model. This technical report provides information on the approach used, model validation tests as well as underlying sources. The aim of this report is to share the results of analysis externally. The rest of this report is structured as follows:

- The ***Model descriptions section*** sets out the two statistical modelling approaches
- The ***Input data and variables section*** sets out the rationale for the input variables chosen in the models as well as the underlying sources
- The ***Model training section*** provides details on the training data used
- The ***Model validation section*** explores the performance in terms of accuracy and reliability of the two models
- The ***Results section*** provides a comparative overview of the results of the two models
- The ***Annex*** includes additional information on the underlying sources and model results.

Model Descriptions

1. Model 1: Global Random Forest

Model 1 uses a single-stage random forest model trained on a broad set of 2,570 jurisdiction-year observations in 297 unique jurisdictions[1], across 57 countries between 2005 and 2014. This set is derived from applying ART TREES HFLD[2] condition to a broad set of jurisdictions globally. It learns from these samples to predict deforestation rates based on a set of 19 covariates. The target variable is a 5 year forward looking, average deforestation rate, derived from Hansen et al.'s Global Forest change data. The model explicitly captures HFLD dynamics and attempts to answer the question "what deforestation rate should we expect in a given jurisdiction, based on these covariates". This prediction then applied to HFLD jurisdictions gives an estimate of counterfactual risk based on similar cases.

A random forest is an ensemble machine learning method, which uses the collective knowledge of many decision trees. Decision trees are simple models that learn to split data into groups based on a series of yes/no questions about the features.

Figure 2 below represents a hypothetical simplified example of a random forest regression model where one of the decision trees in the forest is making decisions based upon the value of the forest area coverage and population density to make a prediction of the deforestation rate.

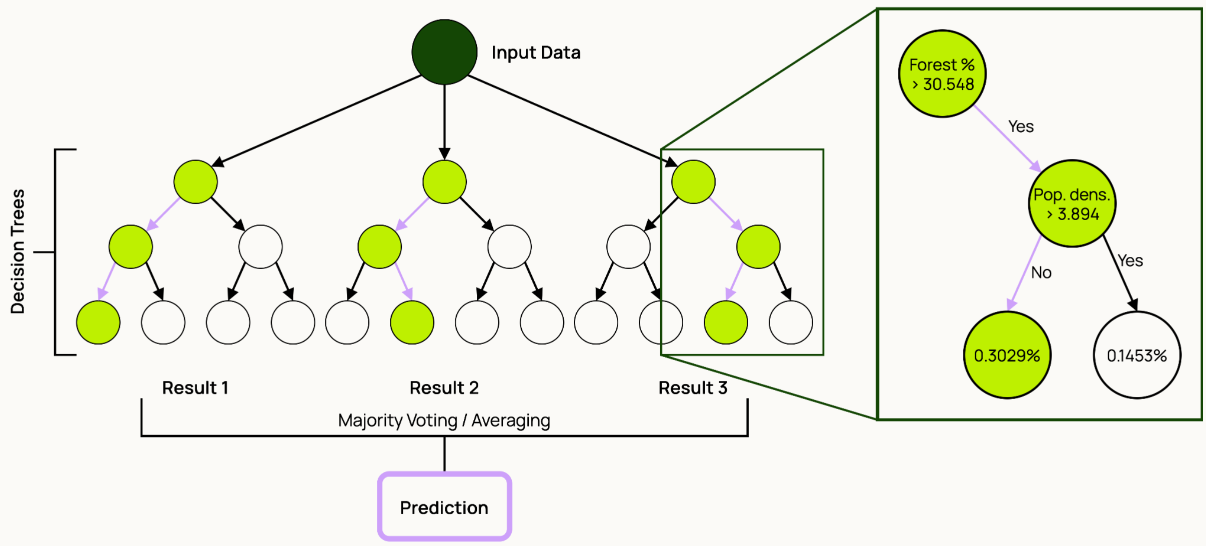

Figure 2. Random forest regression model for predicting deforestation rates
Source: Sylvera, 2026

2. Model 2: Two-Step (Logistic Regression and Random Forest)

Model 2 builds on the methods presented in Teo et al. (2024) paper that estimated future deforestation rates for HFLD jurisdictions using a two-step approach. The logistic step was trained on 2,580 jurisdiction-year observations drawn from 297 unique HFLD jurisdictions, tracking their transitions across

[1] "Jurisdiction" in this study refers to ADM1 level boundaries, the level below national boundaries. Examples include states or provinces.
[2] The ART definition for a High Forest Low Deforestation (HFLD) jurisdiction is a HFLD score of more than 0.5 maintained over a 5 year period. The HFLD score is given by (0.5% - Defor. Rate) + ((Forest Coverage - 50%) / 100). For a deforestation rate of 0.01 and a forest coverage of 56%, the HFLD score would be (0.5 - 0.01) + (0.56 - 0.5) = 0.49 + 0.06 = 0.55.

each year from 2005 to 2014. The Random Forest step was trained only on the transition years of jurisdictions which lost HFLD status which resulted in 160 observations. This two step model first applies a lasso logistic regression for predicting the likelihood that a HFLD jurisdiction will lose its status in the subsequent 5 years (i.e. will 'transition' out of HFLD status).

Logistic regression is a form of supervised machine learning algorithm that describes the process of estimating the relationship between a series of independent variables (in this case, environmental and socio-economic factors) and a given outcome (likelihood of an HFLD transition).

Training a logistic regression model works by estimating the coefficients (parameters) that describe how each input variable contributes to the outcome. It takes inputs, combines them in a linear way, and then uses a logistic (S-shaped) function to turn that into a probability between 0 and 1. Once these coefficients are fitted, they are applied to new data using Eq. (1) to calculate the probability of the outcome occurring. This probability can then be used to classify new observations as either likely or unlikely to transition.

$$\textbf{(1)} \qquad log\left(\frac{p(X)}{1-p(X)}\right) = \beta_0 + \beta_1 X,$$

where $p(X)/1 - p(X)$ is called 'odds' - and is the ratio between the outcome occurring or not. This makes the entire left-hand-side the 'log-odds' function. Solving this equation for $p(X)$ would return a sigmoid function, hence why the output of a logistic regression model is a probability between 0 and 1.

Example: Suppose the fitted coefficients are $\beta 0=-2$ and $\beta 1=0.5$. If $X=4$ then:

$$log\left(\frac{p(X)}{1-p(X)}\right) = 2 + 0.5(4) = 0$$

Converting back from log-odds gives:

$$p(X) = \frac{1}{1+e^{-0}} = 0.5$$

In this example, the model predicts a 50% probability of the outcome occurring when X = 4.

As an alternative to the stepwise regression employed by Teo et al. (2024), Sylvera's model applies LASSO (Least Absolute Shrinkage and Selection Operator) regularization to the logistic regression. LASSO regularization works by adding a penalty term to the model's loss function that is proportional to the absolute value of the coefficients, as shown in Eq. (2):

(2) minimize { -log-likelihood + λ Σ|β▒| }

where λ controls the strength of the penalty. As λ increases, more coefficients are shrunk toward zero and eventually removed from the model entirely, performing automatic variable selection. This makes LASSO particularly well-suited as a replacement for stepwise regression, as it achieves the same goal of filtering irrelevant predictors, but does so within the model fitting process itself rather than as a separate pre-training step.

$$\textbf{(3)}\ logit\left(P(Y_{i,tI} = 1)\right) = \beta_0 + \sum_{k=0}^{4}\sum_{j=1}^{n} \beta_{j,k} X_{i,j,t-k} + \varepsilon_{i,t}$$

The logistic regression step is formally specified in Eq. (3), which models the probability that HFLD jurisdiction *i* transitions to non-HFLD status within a five-year forward window from year t. Here, $P\ Yi,t = 1$ is the probability that HFLD jurisdiction *i* transitioned to non-HFLD between years $t+1$ and $t+5$, for $t = 2005$ *to 2014*. $X_{i,j,t-k}$ denotes the value of predictor variable j for jurisdiction i at year $t-k$, and $\beta_{j,k}$ is the corresponding coefficient for variable j at lag k. β_0 is the model intercept and $\varepsilon_{i,t}$ is the error term.

To select the optimal λ, Ridge (α=0), LASSO (α=1), and Elastic Net (α=0.5) regularization were each evaluated using 10-fold cross-validation minimizing binomial deviance across all five lag-4 training datasets. LASSO produced the best performance and was selected, with the optimal λ identified as 0.000675. This is a notably low value, indicating that the penalty applied to the coefficients was mild — meaning most predictors were retained and the model behaves similarly to standard logistic regression, with only modest variable selection occurring. The classification threshold was set at 0.5, such that jurisdictions with a predicted transition probability exceeding 50% are classified as likely to transition out of HFLD status within the subsequent five years.

Lagged variables are incorporated in the model to capture the idea that deforestation risk in a given year is not solely determined by present conditions, but is also shaped by trends in prior years. For example, sustained declines in HDI or rising population density over several preceding years may be more informative predictors of an imminent HFLD transition than a single year's observation alone. To determine the optimal lag depth, five model configurations were evaluated, incorporating between zero and four years of lagged predictors. A lag of four years was selected based on consistently strong out-of-sample performance relative to lower lag depths, whilst retaining sufficient sample size for stable model estimation (see section *Model Validation*).

In Sylvera's model, the training data for the logistic regression step is the Hansen forest loss data and covariates (see next section) for the HFLD jurisdictions as of 2005. These HFLD jurisdictions were then assigned a binary label according to whether they transitioned out of HFLD status (1), or not (0), at some point between 2006 and 2014. The test set comprises HFLD transitions observed between 2015 and 2020.

A random forest regressor is then used to predict the 5-year forward average deforestation rate for jurisdictions that have lost HFLD status. This model was trained on observations between 2006 and 2014 where a jurisdiction transitioned out of HFLD status in that year.

The transition likelihood (as predicted via logistic regression) and the deforestation rates (as predicted via random forest) are then combined in Eq. (4) to produce a probabilistic weighted average expected deforestation rate for each HFLD jurisdiction as of 2020:

$$\textbf{(4)}\ r_{exp,i} = \left(\frac{5 - yrs_i}{5}\right)(p_i \cdot r_{h,i} + (1 - p_i) \cdot r_{ld,i}) + \left(\frac{yrs_i}{5}\right) r_{ld,i}$$

where

- $r_{exp,i}$ is the expected deforestation rate for the i^{th} HFLD estimated as the weighted predicted average deforestation rate between *t+1* and *t+5* for jurisdiction *i*
- p_i is the estimated probability of jurisdiction *i* losing its HFLD status at some point between *t+1 and* t+5 as predicted by the logistic regression in Step 1a HFLD transition
- $r_{h,i}$ is the 5-year forward average deforestation rate for jurisdiction *i*, conditional on losing HFLD status, as predicted by the random forest model
- $r_{ld,i}$ is the 5-year historical average deforestation rate for jurisdiction *i*, used as the expected rate conditional on maintaining HFLD status and

- yrs_i is the average number of years that jurisdiction *i* individually remained in HFLD status over the 2005–2014 training period - higher values indicate greater forest stability and weight the prediction more toward $r_{ld,i}$

The final rate produced by Equation 4 is a probability-weighted forecast that combines modeled probability outcomes and historic rates to estimate what is likely to happen on average each year over the scenario period. It is normalized by each jurisdiction's historic HFLD tenure.

If the jurisdiction-year pairing is HFLD, the weighted formula is used to produce an expected deforestation over the subsequent 5 years. If the jurisdiction-year is non-HFLD, $r_{ld,i}$ (the observed rate) is used. This approach is taken to produce country-level estimates where some jurisdictions do not have HFLD status (and therefore the above formula cannot be used) and some jurisdictions have HFLD status (and therefore the model outputs can be used).

This approach allows for country-level predictions, even when some jurisdictions are individually ineligible for HFLD status but qualify when results are aggregated at the national level.

Input Data and Variables

There is a wide array of possible deforestation drivers, and not all are included in this model. Sylvera selected a number of key drivers to include in the model based on a literature review including the work by Geist et al, Grinand et al, Jaffe et al, and Cabral et al.

Table 1 below outlines the input datasets used in the models, including their spatial resolution, the geographic level they represent, and what timeframe they represent. "ADM1" refers to subnational units such as states or provinces, while "ADM0" indicates national-level (country) data. A list of data sources can be found in the Appendix. Jurisdictions used FAO global boundaries at ADM0 and ADM1 levels.

The forest loss data used in this model is sourced from Hansen et al.'s Global Forest Change data. This dataset captures forest cover change at a resolution of 30 meters per pixel, enabling accurate and consistent deforestation estimates particularly at sub-national scales. Hansen's data has been shown to have strong global accuracy (approximately 88%) and is updated annually. As an open-source dataset, it is widely used across both the scientific and policy communities. While country-specific methods for detecting deforestation can yield more locally accurate results, they are tailored to particular ecosystems, or administrative contexts, and aren't available for global application. A global product like Hansen et al. provides a consistent data source across all jurisdictions covered in this study. It should be noted that the quality of detections in this product increased after 2015, due to incorporation of Landsat 8, data availability, and improvements to the algorithm.

Table 1. Input variables to models

Variable	Spatial resolution	Area represented	Temporal resolution	Method
Elevation (m asl)	90m	ADM1	Static	Extracted
Slope (°)	90m	ADM1	Static	Extracted
Temperature (°C)	1km	ADM1	Annual	Extracted
Precipitation (mm)	1km	ADM1	Annual	Extracted
GDP per capita (US$)	1km	ADM1	Annual	Extracted
Human Development Index	1km	ADM1	Annual	Extracted
Nightlight intensity	1km	ADM1	Annual	Extracted
Population density (people per km2)	1km	ADM1	Annual	Extracted

Variable	Spatial resolution	Area represented	Temporal resolution	Method
Mining area (%)	Spatially explicit polygons	ADM1	Static for period 2000-2017 - assumes that this is representative of areas with accessible mineral deposits, and that mining infrastructure remains long term	Extracted
Tree plantation area (%)	30m	ADM1	Annual	Extracted
Forest cover (%)	30m	ADM1	Annual	Derived from Hansen global forest extent in 2000, and loss in ha for subsequent years
Rural population (%)	NA	ADM0	Annual	Extracted
Road density	8km	ADM1	Static for 2018 - year on year data is unavailable, and while new roads can be built, this broadly captures connectivity by road at the ADM1 level	Extracted
Patch density[3]	300m	ADM1	Annual	Calculated by annual forest cover rasters using the Landscape Metrics package in R
Mean patch size[4]	300m	ADM1	Annual	Calculated by annual forest cover rasters using the Landscape Metrics package in R
Edge density[5]	300m	ADM1	Annual	Calculated by annual forest cover rasters using the Landscape Metrics package in R
Gold price USD (avg. spot)	NA	Global	Annual	Extracted
Lithium price USD (avg. spot)	NA	Global	Annual	Extracted
Cobalt price USD (avg. spot)	NA	Global	Annual	Extracted

[3] Patch density - the number of distinct forest patches per unit area, indicating how fragmented the landscape is.
[4] Mean patch size - the average size of individual forest patches, with smaller values suggesting greater fragmentation.
[5] Edge density - the total length of forest-non-forest boundaries per unit area, reflecting the amount of edge habitat created by deforestation.

Model Training

The following section details metrics on the training data used for the respective models.

Model 1: Global Random Forest

The training set comprises 2,570 HFLD jurisdiction-year observations spanning 2005–2014, with the dependent variable being the five-year forward-looking mean deforestation rate (t+1 to t+5). The test set comprises 1,016 HFLD jurisdiction-year observations covering 2015–2020, used to evaluate out-of-sample predictive performance prior to generating 2020 predictions.

Table 2. Model 1 (Single stage Random Forest) Training details

Set	Years	Unit of Observation	Sample Size (n)	Independent variables
Training	2005-2014	HFLD jurisdiction - year	2,570	• Elevation (m above sea level) • Slope (degrees) • Mean annual temperature (°C) • Mean annual precipitation (mm) • Forest cover (%) • Tree plantation area (%) • Patch density • Mean patch size • Edge density • GDP per capita ($USD) • Human Development Index (HDI) • Population density (people per km²) • Rural population (% of total population) • Nighttime light intensity • Road density • Mining area (%) • Gold price (average annual spot price, USD) • Cobalt price (average annual spot price, USD) • Lithium price (average annual spot price, USD)
Test	2015-2020	HFLD jurisdiction - year	1,016	

Source: Sylvera, 2026

Model 2: Two-Step (Logistic Regression and Random Forest)

Logistic Regression

The training set comprises 2,580 jurisdiction-year observations monitoring HFLD transitions across each year from 2005 to 2014, drawn from 297 unique HFLD jurisdictions. Of these, 133 (44.8%) transitioned out of HFLD status within a subsequent five-year window and 164 (55.2%) retained their status throughout, with SMOTE applied during training to address this class imbalance.. The test set comprises 207 HFLD jurisdictions observed as of 2015, used to evaluate out-of-sample predictive performance prior to generating 2020 predictions.

Table 3. Logistic regression training details

Set	Years	Unit of Observation	Sample Size (n)	Dependent variable	Independent variables
Training	2005-2014	HFLD jurisdiction - year (time *t*)	Jurisdiction-year observations across 2005–2014 (n = 2,580 before SMOTE[6]; n = 4,114 after SMOTE), drawn from 297 unique HFLD jurisdictions. Of these jurisdictions, 133 (44.8%) transitioned out of HFLD status at least once within a subsequent five-year window, while 164 (55.2%) retained HFLD status throughout.	Binary outcome equal to 1 if an HFLD jurisdiction transitions to non-HFLD status during any year in t+1–t+5, and 0 otherwise	• Elevation (m above sea level) • Slope (degrees) • Mean annual temperature (°C) • Mean annual precipitation (mm) • Forest cover (%) • Tree plantation area (%) • Patch density • Mean patch size • Edge density • GDP per capita ($USD) • Human Development Index (HDI) • Population density (people per km²) • Rural population (% of total population) • Nighttime light intensity • Road density • Mining area (%) Note** commodity price data was excluded after testing the logistic regression mode, as they had significant drift and weak predictive power
Test	2015-2020	HFLD jurisdiction - year (time *t*)	HFLD jurisdiction-year observations as of 2015 (n = 207).		

Source: Sylvera, 2026

[6]SMOTE is a technique that creates artificial examples of rare events by blending characteristics of real rare events, helping the model learn patterns from underrepresented cases when one outcome is much less common than another.

Random Forest

The training set comprises 160 jurisdiction-year observations spanning 2006–2014, where each observation corresponds to the exact year a jurisdiction lost HFLD status. Rather than tracking all HFLD jurisdictions, this dataset is intentionally restricted to transition events only, capturing the covariates and associated five-year forward deforestation rates specific to forests at the point of losing HFLD status. The test set comprises just 10 observations corresponding to jurisdictions that lost HFLD status in 2015, reflecting the relatively rare nature of transition events in any given year. This focused training approach ensures the random forest learns deforestation patterns specific to transitioning forests, which inform the high deforestation rate component of the composite model.

Table 4. Two-Step Logistic Regression and Random Forest - Step 2 Random Forest Training Details

<table>
<tr><th>Set</th><th>Years</th><th>Unit of Observation</th><th>Sample Size (n)</th><th>Dependent variable</th><th>Independent variables</th></tr>
<tr><td>Training</td><td>2006-2014</td><td>Jurisdiction-year observations indexed at baseline time t, corresponding to an HFLD loss event</td><td>160</td><td rowspan="2">Five-year forward-looking mean deforestation rate (t+1 to t+5)</td><td rowspan="2">Elevation (m above sea level)
Slope (degrees)
Mean annual temperature (°C)
Mean annual precipitation (mm)
Forest cover (%)
Tree plantation area (%)
Patch density
Mean patch size
Edge density
GDP per capita ($USD)
Human Development Index (HDI)
Population density (people per km^2)
Rural population (% of total population)
Nighttime light intensity
Road density
Mining area (%)
Gold price (average annual spot price, USD)
Cobalt price (average annual spot price, USD)
Lithium price (average annual spot price, USD)</td></tr>
<tr><td>Test</td><td>2015</td><td>Jurisdiction-year observations indexed at baseline time t, corresponding to an HFLD loss event</td><td>10</td></tr>
</table>

Source: Sylvera, 2026

Model Validation

The following section details the statistical performance metrics of Models 1 and 2 based on their training data as well as independent data that they had not 'seen' during the training process. Since Model 1 and the random forest element of Model 2 are the same technique, just applied to different input data, the metrics used are identical. However, the logistic regression algorithm used by Model 2 - despite its name - is a classification technique. For reference, a classification model predicts a class, whereas a regression model (such as random forest) predicts a non-discrete value. Therefore, the validation section for the logistic regression step reports different metrics to the random forest models. The results of the two models, i.e. the predicted values across a subset of jurisdictions and countries and key takeaways are summarized in the *Results* section.

Accuracy metrics used in assessing final results

For the purposes of this sub-section the following accuracy metrics are presented to assess how far predictions are from actual values.

Explained Variance: the proportion of variability in observed outcomes accounted for by the model. Similar to R^2 but does not penalise systematic bias as it accounts for prediction bias. The percentage term indicates prediction spread relative to observed variance.

RMSE (Root Mean Squared Error): the average magnitude of prediction errors, with larger errors penalised more heavily. Expressed in the same units as the dependent variable (deforestation rate).

MAE (Mean Absolute Error): the average absolute difference between predicted and observed values. Less sensitive to large errors than RMSE.

MAPE (Mean Absolute Percentage Error): the average percentage difference between predicted and observed values. Can provide a simple and interpretable metric of prediction accuracy with lower values reflecting better performance. Can be inflated when observed values are close to zero.

Observed Variance Captured: the ratio of the spread of model predictions to the spread of observed values. A value below 100% indicates the model is compressing predictions toward the mean, a known characteristic of ensemble methods such as random forests.

Model 1: Global Random Forest

The Global Random Forest model was evaluated on a held-out test set of 1,016 HFLD jurisdiction-year observations covering 2015–2020. Table 5 below provides a summary of the key performance metrics for Model 1. The model achieved an explained variance of 0.775, indicating it captures the majority of variation in five-year forward deforestation rates. The RMSE was 0.121 and MAE was 0.061, with minimal directional bias (-0.007), suggesting the model neither consistently over- nor underpredicts. The MAPE of 74% is elevated but expected, as the bulk of observations are concentrated at very low deforestation rates where small absolute errors translate into large percentage differences. Model 1 predictions captured 62% of observed variance, reflecting the tendency of random forest models to compress predictions toward the mean, which is most pronounced at higher deforestation rates driven by episodic events not captured by the model's covariates.

Table 5. Model 1 Random Forest Performance metrics

	MAE	RMSE	% above	% below	MAPE	Explained variance	n test
Model 1	0.061	0.121	51%	49%	74%	0.775 (62%)	1,016

Source: Sylvera, 2026

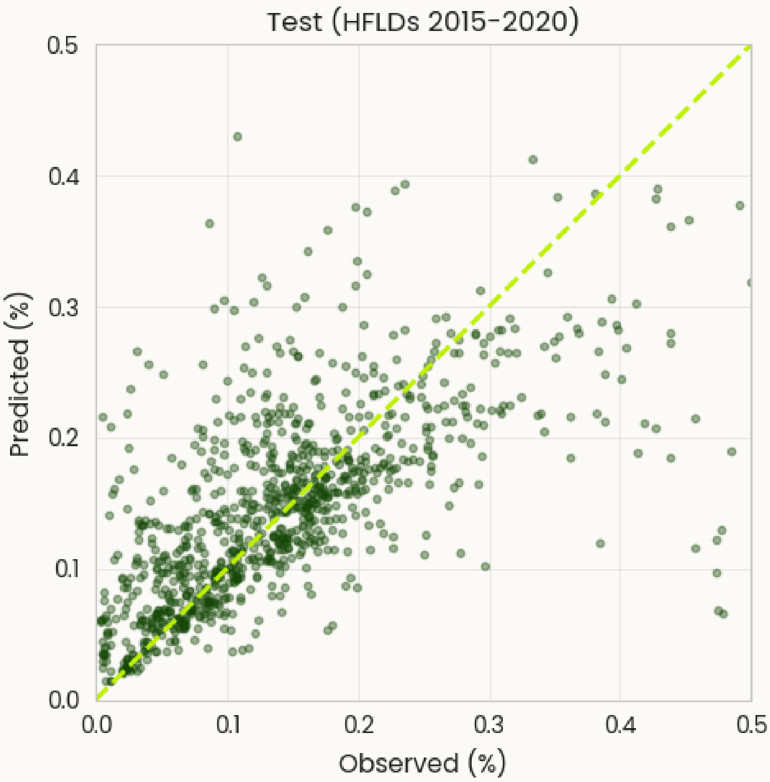

Figure 3. Model 1 Predicted deforestation rates versus observed deforestation rates for test set
Source: Sylvera, 2026

The distribution of Model 1 predictions is heavily right-skewed (Figure 4), consistent with the low deforestation rates characteristic of HFLD jurisdictions. Approximately 70% of predictions fall below 0.2%, with progressively fewer jurisdictions predicted at higher rates and only a small proportion exceeding 0.3%. Very few predictions exceed 0.5% — the maximum deforestation rate permitted for HFLD status under the ART TREES standard — reflecting the model's appropriate concentration of predictions within the low-deforestation range that defines this class of jurisdiction.

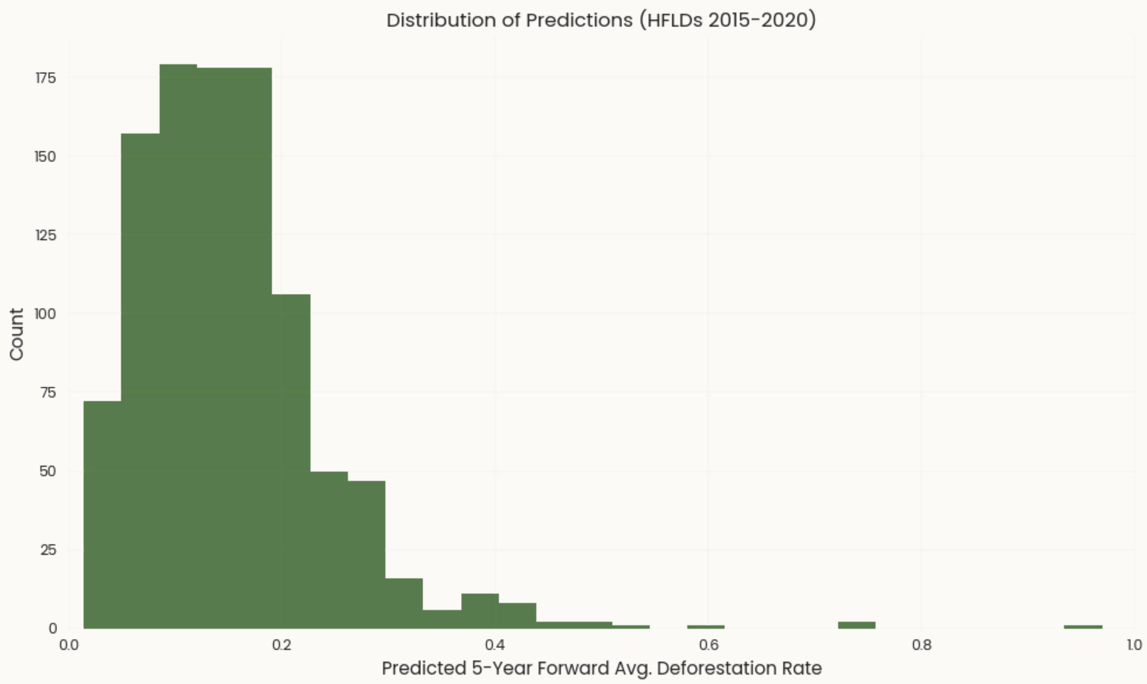

Figure 4. Model 1 Distribution of predicted deforestation rates histogram
Source: Sylvera, 20256

In the following figure (Figure 5) the top 20 variables which influenced the model are provided. A full table of predictive variables that input to the model is provided in the Section *Inputs Data and Variables* and the *Annex*.

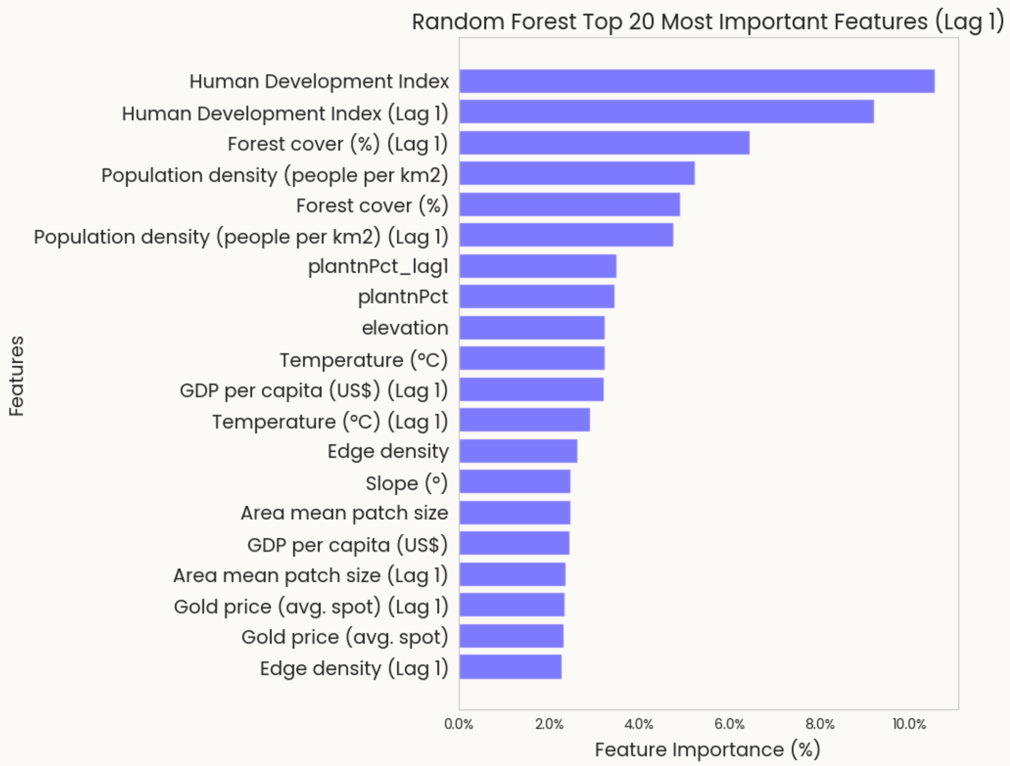

Figure 5. Model 1 variable importance - top 20
Source: Sylvera, 2026

Model 2: Two-Step (Logistic Regression and Random Forest)

This section outlines various outputs in assessing the logistic regression, and the final composite model results using the probabilities with random forest results.

The following metrics are used in describing logistic regression results

AUC (Area Under the ROC Curve): measures the model's ability to distinguish between jurisdictions that will transition out of HFLD status and those that will not. A value of 1.0 indicates perfect discrimination, while 0.5 indicates no better than random chance.

AIC (Akaike Information Criterion): a measure of model fit that penalises complexity, used to compare model configurations. Lower values indicate a better balance between fit and parsimony.

BIC (Bayesian Information Criterion): similar to AIC but applies a stronger penalty for additional parameters. Lower values are preferred.

RMSE (Root Mean Squared Error): the average magnitude of prediction errors in predicted transition probabilities.

Brier Score: measures the mean squared difference between predicted probabilities and actual binary outcomes. Lower values indicate better-calibrated probability predictions, with 0 being perfect and 1 being worst.

Accuracy: the proportion of jurisdictions correctly classified as transitioning or not transitioning.

Precision: of all jurisdictions predicted to transition, the proportion that actually did. High precision indicates few false positives.

Recall: of all jurisdictions that actually transitioned, the proportion correctly identified by the model. High recall indicates few false negatives.

F1 Score: the harmonic mean of precision and recall, balancing the trade-off between the two.

Specificity: of all jurisdictions that did not transition, the proportion correctly identified as stable. High specificity indicates few false positives among non-transitioning jurisdictions.

Logistic Regression – testing optimal lag

The results of the lag evaluation are shown in Table 6, representing the average results for each lag under different evaluation windows. Lag 4 was selected as the optimal model configuration based on strong out-of-sample performance; with an average test AUC of 0.816, and Train AUC of a 0.908, indicating good generalisation without excessive overfitting. The gap between training and test AUC is modest, suggesting the model captures genuine predictive signals rather than memorising training data. AIC values were lower than Lag 3, supporting a better balance between model fit and complexity. AIC penalises models for adding additional parameters that do not meaningfully improve predictive performance, meaning a lower AIC indicates that the additional year of lagged variables incorporated in Lag 4 contributed genuine predictive value rather than simply increasing model complexity.

Table 6. Model 2 logistic regression Performance metrics for lags 0-4

Lag	AIC	BIC	Test RMSE	Test Brier	Test Accuracy	Test Precision	Test Recall	Test F1	Test Specificity	Test AUC	Train AUC
0	2,481	2,592	0.478	0.229	0.634	0.447	0.659	0.551	0.624	0.692	0.746
1	1,980	2,161	0.436	0.188	0.72	0.561	0.635	0.591	0.751	0.796	0.857
2	1,807	2,043	0.421	0.179	0.746	0.628	0.662	0.641	0.787	0.808	0.889
3	1,715	2,028	0.419	0.178	0.75	0.61	0.677	0.638	0.773	0.811	0.901
4	1,678	2,028	0.418	0.177	0.747	0.628	0.675	0.645	0.779	0.816	0.908

Source: Sylvera, 2026

Regularization selection

Figure 6 shows the cross-validation error curves for all three regularisation methods tested — Ridge, LASSO, and Elastic Net — across a range of log(λ) values. For each method, the optimal λ values are marked by dashed vertical lines: the pink line indicates λ.min (the value minimising cross-validation error) and the blue line indicates λ.1se (a more regularised alternative within one standard error of the minimum). Ridge shows a gradual increase in error as λ increases, reflecting its behaviour of slowly shrinking all coefficients rather than eliminating them. LASSO and Elastic Net both show a sharp increase in error at higher λ values, consistent with their more aggressive coefficient elimination at stronger penalty strengths. LASSO achieved the lowest minimum cross-validation error of the three methods and was therefore selected as the final regularisation approach.

Figure 7 shows the LASSO cross-validation curve in greater detail, with confidence intervals displayed around the error estimates. The curve is flat and stable at low λ values (left of the plot), indicating that a wide range of low penalty strengths produce similarly good predictions. The selected λ.min (blue dashed line) sits well within this stable region at approximately e^{-7}, confirming that the optimal penalty is mild. The sharp rise in error at higher λ values (right of the plot, around e^{-4} marked by the green dashed line representing λ.1se) illustrates the point at which regularisation becomes too aggressive and begins removing predictive variables, degrading model performance.

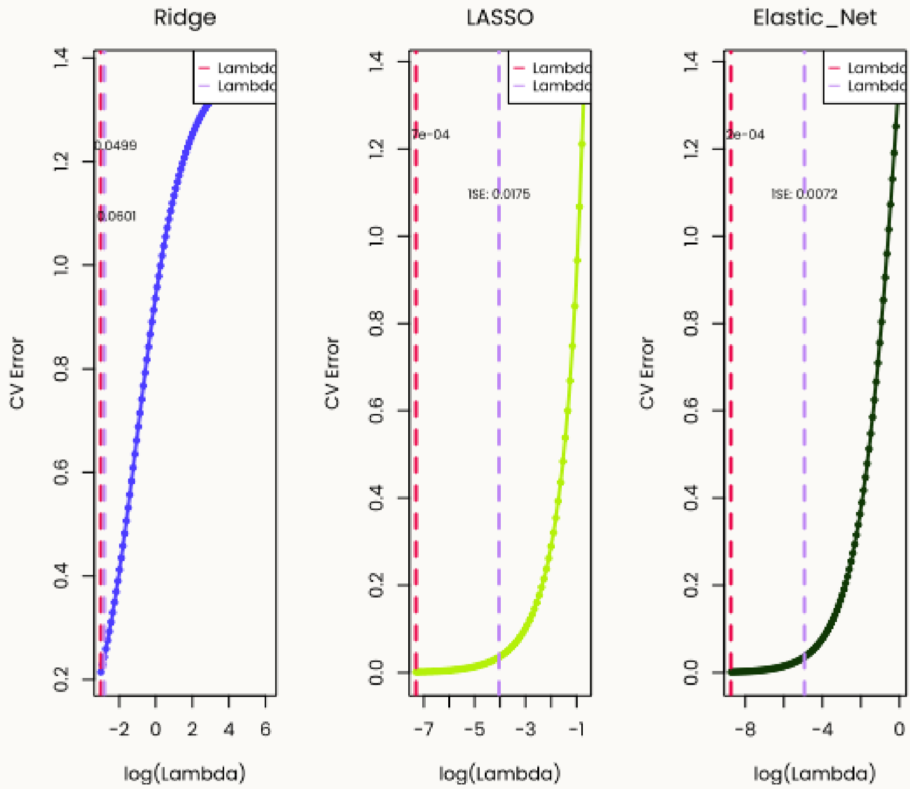

Figure 6. Cross validation error curves for different regularisation methods
Source: Sylvera, 2026

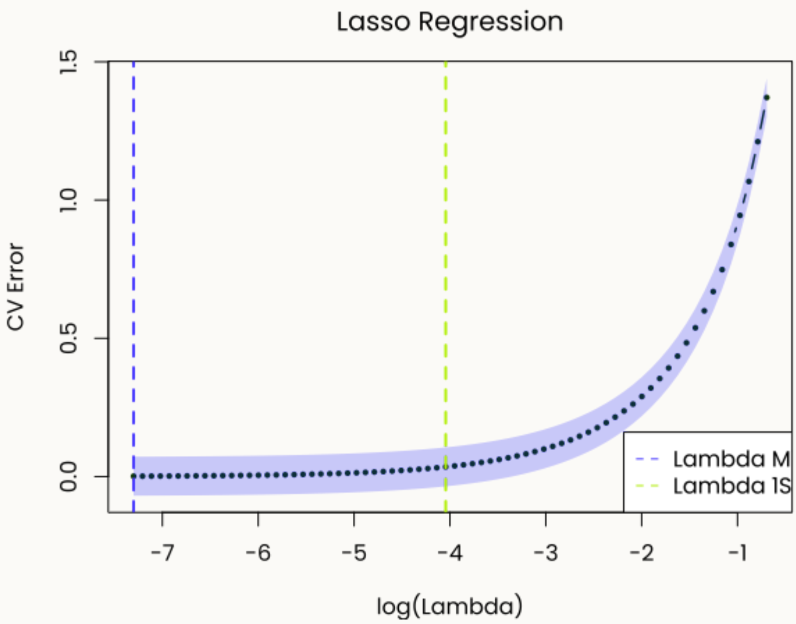

Figure 7. Lasso cross validation curve
Source: Sylvera, 2026

Logistic regression results

Figure 8 shows the ROC curve for the LASSO logistic regression model evaluated on the test set. The curve illustrates the trade-off between correctly identifying transitions (True Positive Rate, y-axis) and incorrectly flagging stable jurisdictions as transitioning (False Positive Rate, x-axis) across all possible classification thresholds. The closer the curve hugs the top-left corner, the better the model's discriminative ability. The model achieves an AUC of 0.814, meaning it correctly ranks approximately 81% of transition versus non-transition pairs. The green point marks the selected classification threshold of 0.5, which sits at a false positive rate of approximately 0.30 and a true positive rate of approximately 0.75, reflecting the model's conservative tendency to prioritise correctly identifying stable jurisdictions over capturing every transition.

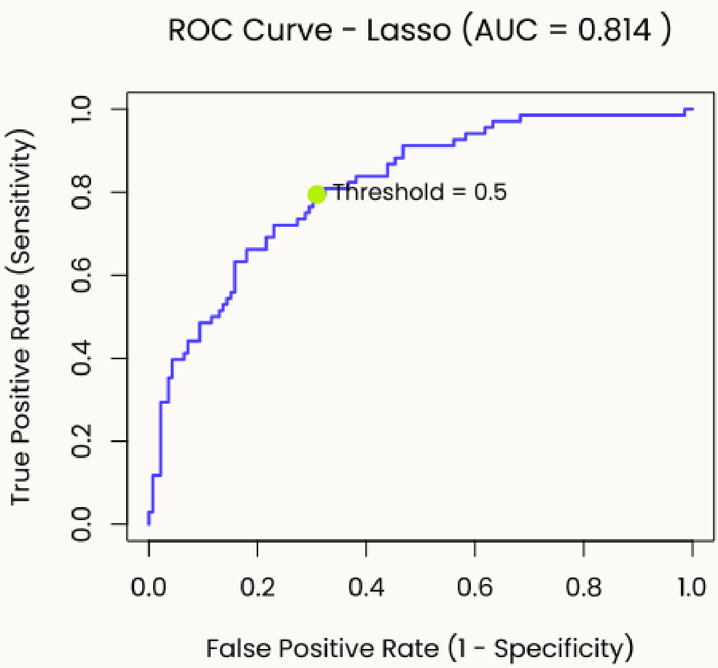

Figure 8. Logistic regression ROC curve
Source: Sylvera, 2026

The full set of performance metrics at the 0.5 threshold are summarised in Table 7 below. The table shows the performance metrics post-regularisation, which is the model used for the composite step. The model correctly classifies 72.5% of observations overall. When it predicts a transition, it is correct 79% of the time (precision), and it correctly identifies 87% of stable jurisdictions (specificity). The model captures 56% of true transitions (sensitivity/recall), which is the weakest metric, reflecting the inherent difficulty of identifying the minority transition class even after SMOTE balancing. The Brier Score of 0.182 confirms that the predicted probabilities are well-calibrated, falling below the 0.2 threshold generally considered indicative of good probability estimation.

Table 7. Model 2 regularised logistic regression Performance metrics (post-regularisation)

Metric	Value	Interpretation
AUC	0.814	Model correctly ranks ~81% of transition vs. non-transition pairs, which is the area under the purple curve on the plot
Accuracy	0.7246	Overall share of correctly classified transitions
Precision	0.7941	When the model predicts a transition its right 79% of the time
Sensitivity (recall)	0.5567	Captures ~56% of true transitions
Specificity	0.8727	Correctly identifies ~87% of non-transitions
Brier Score	0.1824	Measures calibration of probabilities — lower is better; this is good (<0.2)

Source: Sylvera, 2026

Model 2 Random Forest

Figure 9 shows the parity plot for the Random Forest model evaluated on the 10 jurisdictions that lost HFLD status in 2015, where the dashed yellow line represents perfect prediction. Predictions are clustered in the 0–0.5% range, consistent with the low deforestation rates observed even among transitioning jurisdictions. Most points sit above the parity line, indicating a mild tendency to overpredict, though given the small test set size of 10 observations these results should be interpreted cautiously. Figure 10 shows the top 20 variables influencing the model.

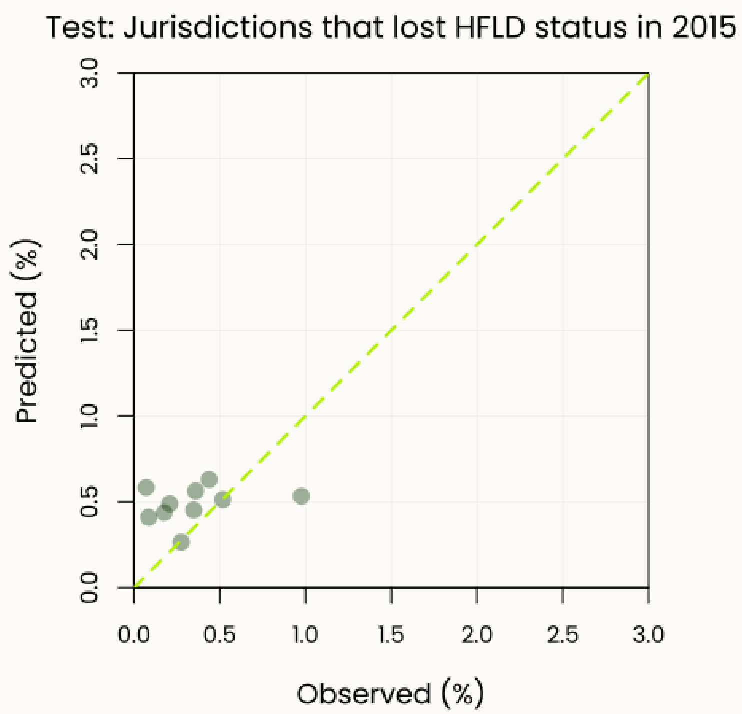

Figure 9. Model 2 random forest - predicted deforestation rates versus observed deforestation rates for test set
Source: Sylvera, 2026

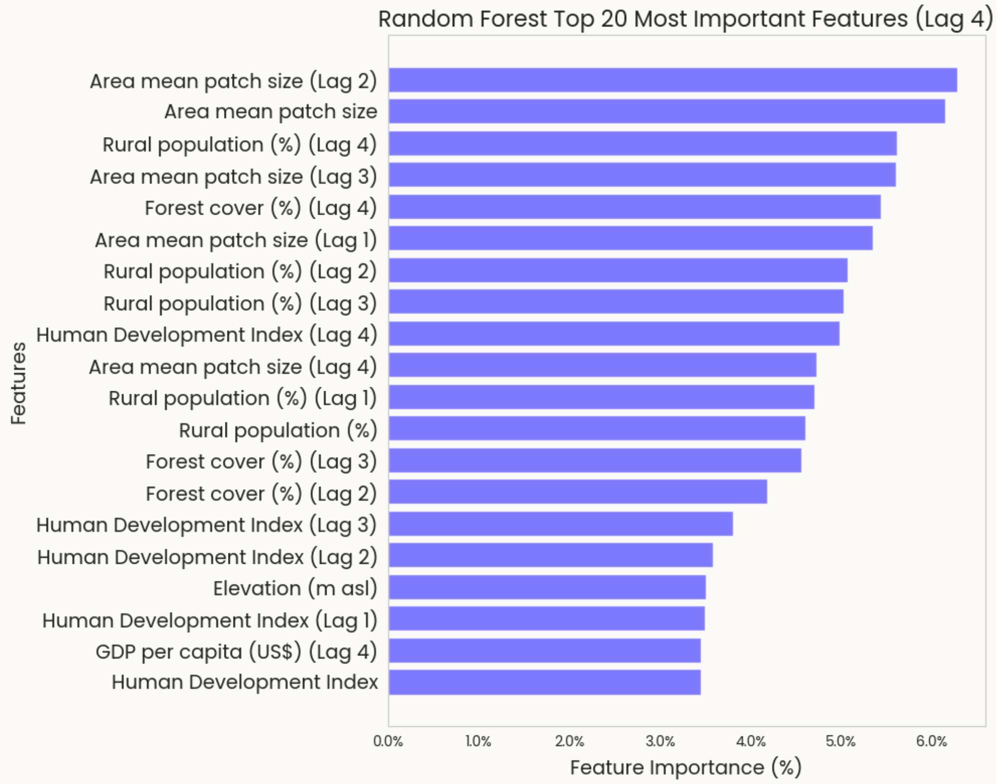

Figure 10. Model 1 random forest variable importance - top 20
Source: Sylvera, 2026

Two-step model - results of final composite

The composite model was evaluated on the same test set as Model 1 with Table 8 providing a summary of the performance metrics. The model achieved an explained variance of 0.055, capturing 12% of observed variance, indicating that predictions are compressed toward the mean. The MAE of 0.076 and RMSE of 0.251 reflect this conservative tendency, with 59% of predictions falling below observed values confirming a systematic underprediction bias.

As observed in Figure 11, the model performs better for jurisdictions with low deforestation rates below 0.2%, where the majority of HFLD jurisdictions are concentrated, but loses predictive ability at higher observed rates. The MAPE of 54% reflects the concentration of predictions in the low deforestation range characteristic of HFLD jurisdictions. Overall, Model 2's conservative bias makes it more suitable for identifying low-risk jurisdictions than for flagging high-deforestation scenarios.

Table 8. Model 2 two-step Performance metrics

	MAE	RMSE	% above	% below	MAPE	Explained variance	n test
Model 1	0.076	0.251	41%	59%	54%	0.055 (12% observed)	1,016

Source: Sylvera, 2026

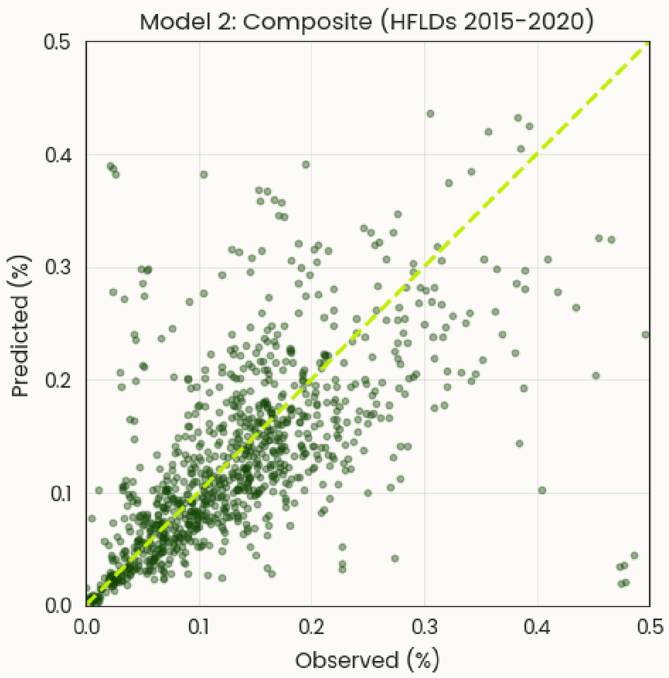

Figure 11. Model 2 composite results - predicted deforestation rates versus observed deforestation rates for test set
Source: Sylvera, 2026

Results and Takeaways

Key Takeaways

Deforestation risk across HFLDs varies. The results show that the number of HFLD jurisdictions has declined over time and that those that have lost their status show increasing deforestation rates. However, for HFLDs that have maintained their status deforestation rates have remained low indicating that a subset of HFLDs will remain at low risk.

Model 1 (Single-stage Random Forest) is best at capturing variability and distribution but Model 2 (Two-step Logistic Regression and Random Forest) can offer a more conservative approach. Although results show that Model 1 is more accurate at predicting deforestation rates for 2015-2020, effectively capturing both variance and extreme values (lower MAE and RMSE values), Model 2 can offer a more conservative approach as predictions are clustered at lower values and can be more suitable for jurisdictions with low rates.

JREDD historical average is unable to capture deforestation spikes. Data shows that the classic JREDD baseline approach - historical average - is unable to capture observed spikes in deforestation and, therefore, cannot predict high risk to HFLDs that will not have been observed in the past.

Model 2's logistic regression step results can be used as a 'screening' tool to indicate risk. Model 2 logistic results can be used as an initial 'screening' tool to identify HFLDs that are at risk of transition, i.e. the logistic step identifies HFLD jurisdictions that are at risk of transitioning out of status at a certain point in time.

Country and spatially-explicit models for improved accuracy. The variation in performance across the models indicates that country-specific models may be necessary for more accurate predictions of deforestation rates of individual countries and / or jurisdictions.

Modelling results

Figure 12 shows the average predicted deforestation rates for HFLD jurisdictions as of 2020, sorted by ascending historical average, alongside observed rates (2021–2024) and the historical average (2016–2020). The chart illustrates how each method tracks observed deforestation across the full distribution of jurisdictions.

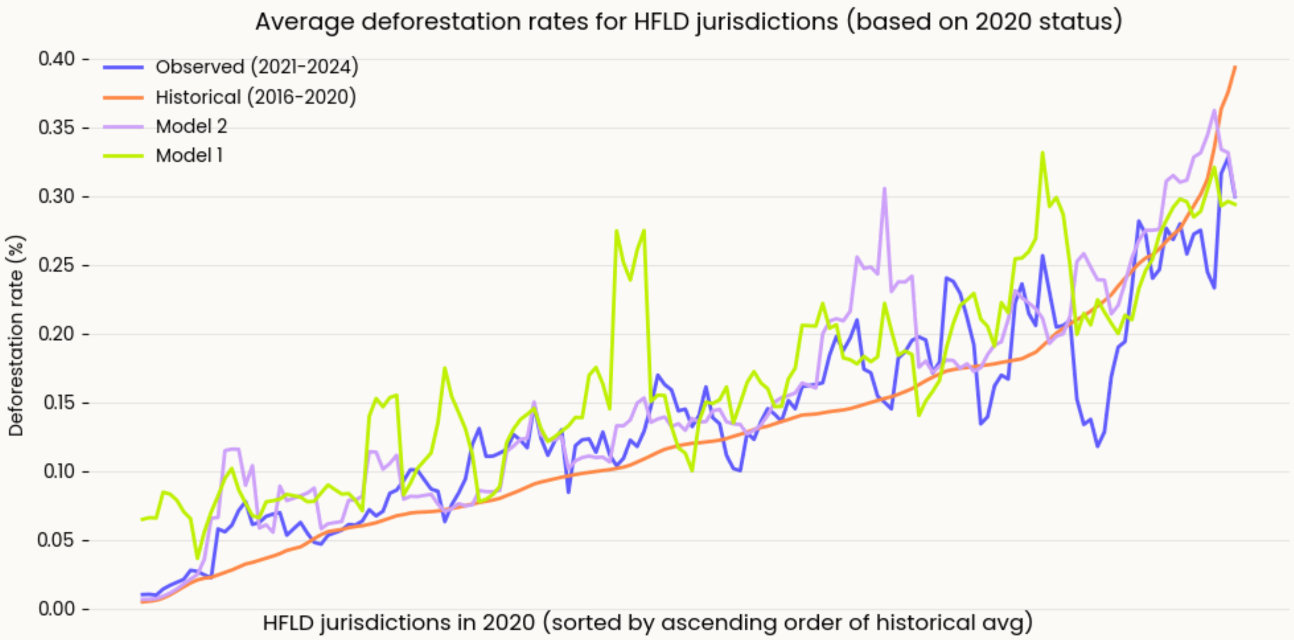

Figure 12. Distribution of values per model for HFLD's as of 2020
Source: Sylvera, 2026

The results across both models and the historical baseline are summarised in the table below. Model 1 achieves the strongest performance on accuracy metrics, with the lowest MAE (0.061) and RMSE (0.121), and captures 62% of observed variance — demonstrating its ability to track higher deforestation states that other methods miss. Its predictions are broadly balanced, with 51% above and 49% below observed values, and a median predicted rate of 0.140% against an observed median of 0.130%.

Model 2 produces narrower, more conservative predictions clustered at lower deforestation rates, reflected in its lower explained variance (12% of observed) and higher RMSE (0.251). However, it achieves a better MAPE (54% vs 74%), driven by its concentration of predictions in the low-deforestation range where percentage errors are smaller. Its median prediction of 0.120% sits slightly below the observed median of 0.130%, consistent with its systematic underprediction tendency.

The historical average performs similarly to Model 2 on MAE (0.077) but captures only 9% of observed variance, confirming that static historical rates are unable to account for jurisdictions where deforestation risk is increasing beyond what past rates would suggest — a key limitation for REDD+ baseline construction in an HFLD context.

Table 9. Performance metrics for Models 1 and, 2 and using a historical average to predict the future deforestation risk

Method	MAE	RMSE	% above	% below	MAPE	Explained variance	Median
Model 1	0.061	0.121	51%	49%	74%	0.775 (62% observed)	0.140 (vs 0.130)
Model 2	0.077	0.251	41%	59%	54%	0.055 (12% observed)	0.120 (vs 0.130)
Historical average	0.077	0.257	30%	70%	42%	0.004 (9% observed)	0.102 (vs 0.130)

Source: Sylvera, 2026

[Figure 13](#) compares the distribution of 5-year forward average deforestation rate predictions across all methods against observed values. The width of each violin at any given rate represents the density of predictions at that value, with wider sections indicating more jurisdictions predicted at that rate. The solid horizontal line marks the median, the dashed line the mean, and the vertical bar the interquartile range.

The observed distribution (dark green) shows the widest spread, with a long upper tail reflecting the minority of jurisdictions experiencing significantly elevated deforestation rates. Model 1 (yellow-green) most closely replicates this distribution, capturing both the concentration of values at lower rates and the presence of higher-rate predictions in the upper tail, consistent with its stronger variance capture. Model 2 (purple) shows a narrower distribution with greater density at lower deforestation rates, confirming its tendency to underestimate variability and compress predictions toward the mean. The historical average (light blue) produces the narrowest distribution of all four, with both its mean and median sitting substantially below the observed values — illustrating that static historical rates systematically fail to capture forward-looking deforestation risk, particularly for jurisdictions where rates are trending upward.

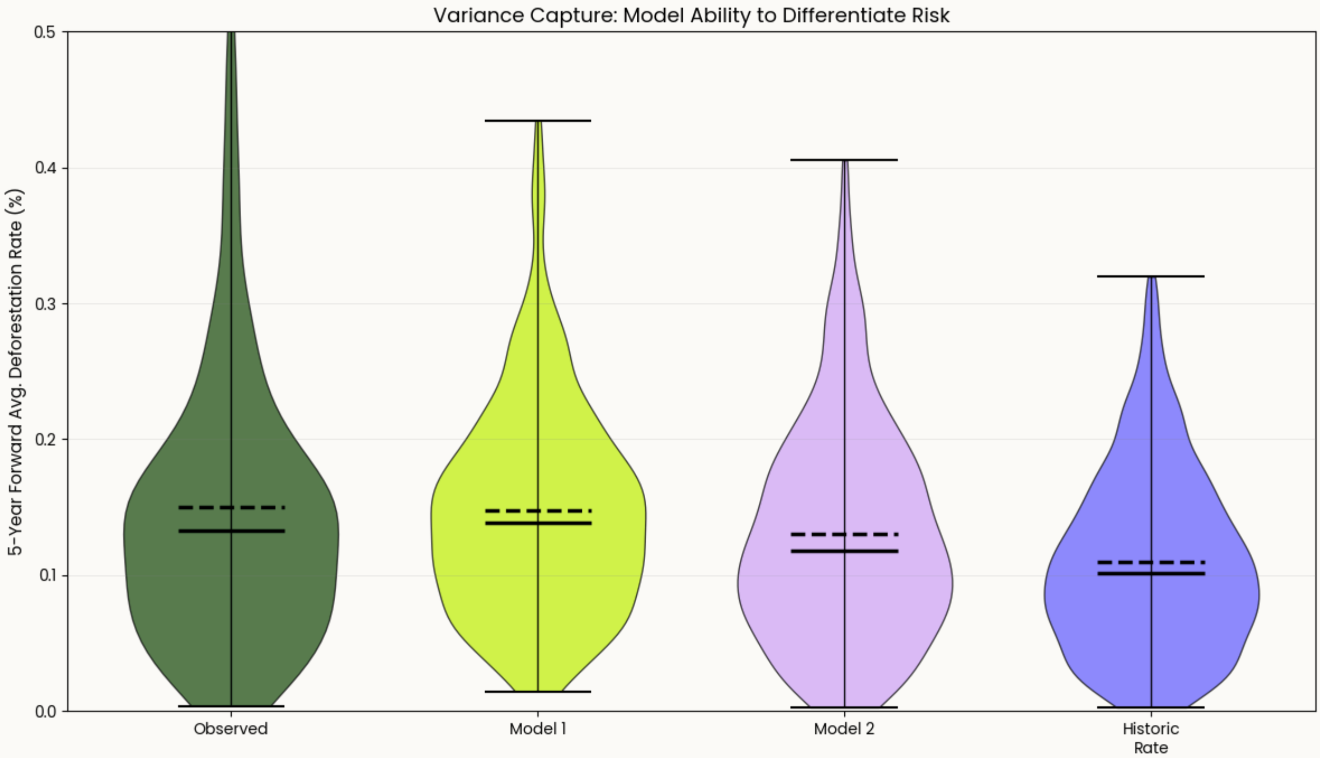

Figure 13. Violin chart showing distribution of deforestation rates
Source: Sylvera, 2026

Transition probabilities as a risk indicator

Beyond their role in the composite deforestation rate forecast, the transition probabilities produced by the LASSO logistic regression step have potential as a standalone screening tool for identifying HFLD jurisdictions at risk of losing their status. When HFLDs do transition out of status, deforestation rates typically accelerate significantly, making early identification of at-risk jurisdictions valuable for both conservation planning and carbon credit risk assessment. At the same time, a persistent cohort of jurisdictions demonstrate stable deforestation rates over time and are unlikely to transition in the near term.

Figure 14 shows the distribution of predicted transition probabilities for the 162 HFLD jurisdictions as of 2020, representing the estimated likelihood that each jurisdiction will lose its HFLD status at some point between 2021 and 2025. Jurisdictions are grouped into three risk categories: low risk (0–0.33, n=43), moderate risk (0.33–0.66, n=70), and high risk (>0.66, n=49). The distribution is notably spread across the full probability range, with no strong concentration at either extreme, suggesting meaningful differentiation in predicted risk across jurisdictions. The largest share of jurisdictions fall in the moderate risk category, while 49 jurisdictions — approximately 30% of the total — are classified as high risk, warranting closer monitoring.

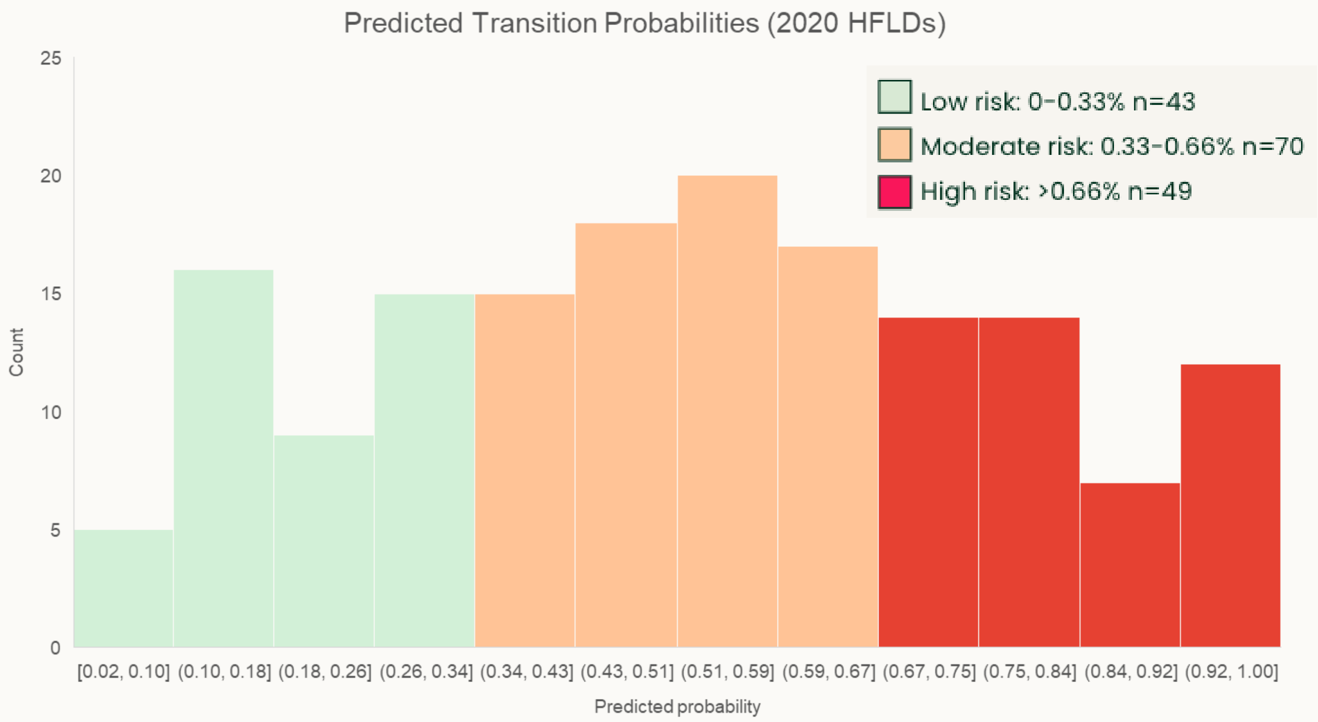

Figure 14. Histogram of transition probabilities with a qualitative risk assigned to them
Source: Sylvera, 2026

Annex

HFLD Definition

For the purposes of this analysis and report we have applied ART TREES definition of HFLD to all jurisdictions in our input dataset.

To qualify as HFLD under ART TREES a jurisdiction must:

- Have >50% forest cover, and
- Maintain an annual deforestation rate <0.5% during the reference period.
- Maintained an HFLD score <0.5 for the preceding 5 years

Only if both conditions are met can the HFLD Score be calculated. The calculation is shown below:

$$HFLD\ Score\ =\ (0.5\ -\ Deforestation\ Rate)\ +\ \frac{Forest\ Cover\ \%-50}{100}$$

When the HFLD Score is greater than 0.5 for each of the five years of the reference period, the jurisdiction is classified as HFLD.

HFLD transition data

Table 10 below provides more detail on the transitions of HFLDs during the study period. It shows the total number of HFLDS, the count and percentage of HFLDs that transitioned in the subsequent 5 years, the mean years of HFLDs remaining HFLD in the subsequent 5 years, as well as the the count and percentage of HFLDs with deforestation rates in the future five years less than the historical average.

Table 10. Details on HFLD transitions

Year (t)	HFLDs as of Year t	HFLDs that transitioned in time t+1 to t+5	Mean years of HFLDS remaining HFLDs from time t+1 to t+5	HFLDs with average deforestation rates in time t+1 to t+5 < historical 5-year deforestation rate
2005	308	86 (27.92%)	4.21	35 (11.36%)
2006	303	79 (26.07%)	4.14	45 (14.85%)
2007	288	74 (25.69%)	4.14	45 (15.62%)
2008	272	73 (26.84%)	4.11	54 (19.85%)
2009	262	78 (29.77%)	4.1	61 (23.28%)
2010	244	66 (27.05%)	4.2	80 (32.79%)
2011	240	82 (34.17%)	3.99	58 (24.17%)
2012	234	86 (36.75%)	3.87	46 (19.66%)
2013	224	80 (35.71%)	3.8	37 (16.52%)

Year (t)	HFLDs as of Year t	HFLDs that transitioned in time t+1 to t+5	Mean years of HFLDS remaining HFLDs from time t+1 to t+5	HFLDs with average deforestation rates in time t+1 to t+5 < historical 5-year deforestation rate
2014	206	63 (30.58%)	3.89	41 (20.0%)
2015	206	68 (33.01%)	3.63	30 (14.56%)
2016	171	34 (19.88%)	4.19	43 (25.15%)
2017	161	26 (16.15%)	4.53	62 (38.51%)
2018	155	28 (18.06%)	4.53	64 (41.29%)
2019	161	42 (26.09%)	4.3	59 (36.65%)
2020	162	33 (20.37%)	3.57	79 (49.69%)

Data sources

Table 11. List of data sources

Input Variable	Source
Elevation (m asl)	Jarvis, A., H.I. Reuter, A. Nelson, E. Guevara. 2008. Hole-filled SRTM for the globe Version 4, available from the CGIAR-CSI SRTM 90m Database: https://srtm.csi.cgiar.org.
Slope (°)	Jarvis, A., H.I. Reuter, A. Nelson, E. Guevara. 2008. Hole-filled SRTM for the globe Version 4, available from the CGIAR-CSI SRTM 90m Database: https://srtm.csi.cgiar.org.
Temperature (°C)	Copernicus Climate Change Service, Climate Data Store, (2023): ERA5 hourly data on single levels from 1940 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: 10.24381/cds.adbb2d47
Precipitation (mm)	Funk, Chris, Pete Peterson, Martin Landsfeld, Diego Pedreros, James Verdin, Shraddhanand Shukla, Gregory Husak, James Rowland, Laura Harrison, Andrew Hoell & Joel Michaelsen. "The climate hazards infrared precipitation with stations-a new environmental record for monitoring extremes". Scientific Data 2, 150066. https://doi:10.1038/sdata.2015.66 2015
GDP per capita (US$)	Kummu, M., Kosonen, M., & Masoumzadeh Sayyar, S. (2025). Data for: Downscaled gridded global dataset for Gross Domestic Product (GDP) per capita at purchasing power parity (PPP) over 1990-2022 [Data set]. Zenodo. https://doi.org/10.5281/zenodo.16741980
Human Development Index	United Nations Development Programme. (n.d.). *Human Development Index*. UNDP Human Development Data Center. https://hdr.undp.org/data-center/human-development-index
Nightlight intensity	Zhao, C., Cao, X., Chen, X. et al. A consistent and corrected nighttime light dataset (CCNL 1992–2013) from DMSP-OLS data. Sci Data 9, 424 (2022). https://doi.org/10.1038/s41597-022-01540-x
Population density (people per km2)	Liu, X., de Sherbinin, A., & Zhan, Y. (2019). Mapping Urban Extent at Large Spatial Scales Using Machine Learning Methods with VIIRS Nighttime Light and MODIS Daytime NDVI Data. Remote Sensing, 11(10), 1247. https://doi.org/10.3390/rs11101247

Mining area (%)	Maus, Victor, Stefan Giljum, Jakob Gutschlhofer, Dieison M. da Silva, Michael Probst, Sidnei LB Gass, Sebastian Luckeneder, Mirko Lieber, and Ian McCallum. "A global-scale data set of mining areas." Scientific Data 7, no. 1 (2020): 1-13. https://gee-community-catalog.org/projects/global_mining/
Tree plantation area (%)	Du, Zhenrong; Yu, Le; Yang, Jianyu; Xu, Yidi; Chen, Bin; Peng, Shushi; et al. (2022). A global map of planting years of plantations. figshare. Dataset. https://doi.org/10.6084/m9.figshare.19070084.v1
Forest cover (%)*	Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini, L., Justice, C. O., & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850-853. https://doi.org/10.1126/science.1244693 Guyana, ART TREES Registry, https://art.apx.com/mymodule/reg/TabDocuments.asp?r=111&ad=Prpt&act=update&type=PRO&aProj=pub&tablename=doc&id1=102
Rural population (%)	World Bank. (n.d.). Rural population (% of total population) [Indicator SP.RUR.TOTL.ZS]. World Bank Open Data. Based on United Nations Population Division's World Urbanization Prospects: 2018 Revision. https://data.worldbank.org/indicator/SP.RUR.TOTL.ZS
Road density	Meijer, J.R., Huijbregts, M.A.J., Schotten, C.G.J. and Schipper, A.M. (2018): Global patterns of current and future road infrastructure. Environmental Research Letters, 13-064006. Data is available at www.globio.info
Patch density	Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini, L., Justice, C. O., & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850-853. https://doi.org/10.1126/science.1244693
Mean patch size	Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini, L., Justice, C. O., & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850-853. https://doi.org/10.1126/science.1244693
Edge density	Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini, L., Justice, C. O., & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. Science, 342(6160), 850-853. https://doi.org/10.1126/science.1244693

Disclaimer

Learn more at sylvera.com

Incentivizing investment in real climate action