MapBiomas Patagonia

Collection 3

Version 1

General coordinators
Ana Eljall

Technical coordinators
Camilo Bagnato

Development
Iván Barberá, Lican Martínez, Fabio Trinco

Team
Daniela Arpigiani, Fernando Salvaré, Camila Mateo, José Daniel Lencinas, Gonzalo Dieguez, Ramon Peña Agrest, Jaime Moyano, Pablo Masera, Cecilia Evequoz, Clara Rey, Boris Díaz, Thomas Kitzberger, Juan Gowda, Leonardo Collado

1. Introduction

1.1. Scope and content of the document

This document describes the theoretical basis, the input data and the step-by-step procedure used to produce the annual land use and land cover (LULC) maps of MapBiomas Patagonia Collection 3. It covers satellite image processing, the feature space used by the classifier, the training-sample strategy, the supervised classification itself and the post-classification filtering and reclassification rules. Collection 3 delivers annual maps for 1985-2025; the classification is actually computed for 1986 and for the continuous series 1998-2025, and the 1986 map is replicated over 1985 and 1987-1997, for which no Landsat imagery of sufficient quality and density is available in the region.

Two components of Collection 3 are documented separately because they do not follow the general pixel-based workflow described here: rainfed forest plantations (mapped with a dedicated object-based methodology, see the companion document "Methodology for Classifying Rainfed Plantations in Patagonia") and the accuracy assessment (see the MapBiomas Argentina validation ATBD). Urban areas and glaciers are produced by national cross-cutting working groups and are integrated into the final product rather than classified within the Patagonia workflow.

1.2. Overview of the MapBiomas initiative

The MapBiomas initiative combines people, algorithms, satellite information and large-scale processing in a methodology designed for the operational production of LULC maps over very large territories. MapBiomas provides an environment in which collaborators from different institutions and disciplines share skills and criteria around a common practice of learning by doing. Thanks to Google Earth Engine (GEE) and open-source technology it is possible to access and process large-scale satellite datasets such as the annual mosaics generated by the MapBiomas network.

MapBiomas collections aim to contribute a fast, reliable, collaborative and low-cost method to generate historical time series of annual land use and land cover maps. All data, maps, code and statistics are openly accessible through the MapBiomas platform. This is possible thanks to (i) Google Earth Engine, which provides the data, the processing algorithms and the cloud infrastructure; (ii) the freely available Landsat time series; and (iii) the MapBiomas collaborative network of organizations and experts. MapBiomas maps are an evolving product, and subsequent collections supersede earlier ones.

1.3. Region of interest and zonification

The Patagonia initiative produces annual LULC maps for southern Argentina. The region of interest extends over approximately 838,000 km² and includes most of the Patagonian (Valdivian temperate and Magellanic subpolar) forests, the Patagonian steppe, the southern portion of the Monte ecoregion and the South Atlantic Islands (Burkart et al., 1999; Olson et al., 2001).

The territory is divided into 13 subregions, or zones (Figure 1). The zonification has two purposes: it reduces spectral confusion by keeping the classifier local to relatively homogeneous ecological contexts, and it splits the computation into tractable units. The zone boundaries were derived by applying the Simple Non-Iterative Clustering (SNIC) segmentation algorithm to a set of ecosystem functional attributes derived from MODIS time series - descriptors of the annual

dynamics of vegetation productivity such as its mean level, its seasonality and the timing of its maximum - and then adjusting the resulting clusters by hand against regional vegetation units (Oyarzábal et al., 2018) and expert knowledge of the region. Functional attributes were preferred over a purely spectral composite because they describe how an ecosystem behaves through the year rather than how it looks on a given date, which is what the zonification is meant to keep homogeneous. Each zone is further subdivided by a hexagonal grid that organizes the visual inspection and sample-collection work of the interpreters: cells are assigned to technicians, and sampling effort is distributed across cells so that the internal heterogeneity of each zone is represented.

Two buffered versions of the zonification are used downstream: a 30 m buffer, to close hairline gaps between adjacent polygons when the per-zone classifications are integrated, and a 5 km buffer, as a fallback so that the integration leaves no holes along zone seams (Section 5.1).

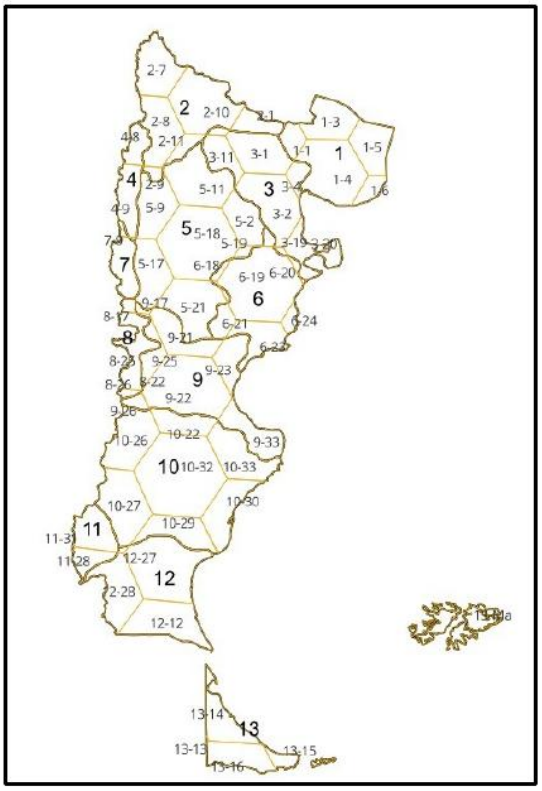

Figure 1. The 13 subregions (zones) of MapBiomas Patagonia Collection 3 and the hexagonal cells used to organize sample collection and map review.

1.4. What changed with respect to Collections 1 and 2

Collection 3 keeps the overall architecture of the previous collections - annual Landsat mosaics, per-zone Random Forest classification trained on temporally stable samples, and a chain of spatial and temporal post-classification filters - but differs from them in several substantive ways:

- Temporal coverage. The series is extended to 2025 (Collection 1 ended in 2022, Collection 2 in 2024).
- No new manual digitization of stable samples. Collections 1 and 2 began with interpreters drawing polygons of temporally invariant classes, from which a

preliminary classification and a stable-area mask were built. In Collection 3 that stage was skipped: stable areas were derived directly from the final, fully post-processed Collection 2 maps (Section 4.3), so production effectively starts at Step 3.

- Expanded legend. From 11 classes in Collection 1 to 14, with Flooded woody vegetation, Peatlands, High-Andean rock and Deserts and badlands separated as distinct classes (Section 4.1).
- Larger feature space. The subset of MapBiomas mosaic variables used by the classifier grew from 31 to 85 variables, and a set of geomorphometric descriptors is computed on the fly and attached to the annual mosaics (Section 3.3).
- Class-probability outputs. The classifier is also run in multi-probability mode and the second and third most probable classes per pixel and year are exported, as evidence for the conceptual post-processing of Step 13.
- Cross-cutting products. Forest plantations are mapped with a dedicated object-based methodology; urban areas and glaciers come from national cross-cutting groups. These classes are excluded from the training set of the general classifier and integrated afterwards (Section 6).
- Independent validation. Collection 3 incorporates a probabilistic accuracy assessment based on an independent stratified reference sample interpreted in Collect Earth (Section 7); Collections 1 and 2 had none.

2. Overview of the methodological process

Figure 2 summarizes the Collection 3 workflow. Annual Landsat mosaics and the associated spectral feature space are produced by the MapBiomas mosaic team and consumed as an input. The Patagonia team defines the legend and the zonification (Steps 1-2 in the figure; in Collection 3 both were inherited, with adjustments, from Collection 2).

The production chain properly begins at Step 3. Areas that kept the same class throughout the Collection 2 series are identified and used as a stable-area mask (Step 3). Random samples are drawn inside that mask, stratified and balanced by the relative area of each class within each zone, and the spectral values of every annual mosaic are extracted at those points (Step 4). A first Random Forest classification - the "stable pre-classification" - is produced per zone and year from those samples alone (Step 4). Interpreters then review the resulting maps zone by zone and hexagon by hexagon and digitize complementary samples wherever the map is wrong; a new Random Forest is trained on the union of stable and complementary samples and re-run for all years (Step 5). This review-resample-reclassify loop is iterated until the interpreters consider that the remaining errors cannot be fixed with more samples.

The 13 per-zone classifications are then integrated into a single Patagonia-wide map, no-data pixels are filled from the temporally nearest valid classification, and a chain of spatial and temporal filters removes noise and stabilizes the time series (Steps 6-12). Finally, conceptual reclassification rules - based on elevation, slope and hand-delimited sectors - correct systematic confusions the statistical filters cannot resolve, and the cross-cutting layers are integrated (Step 13). Accuracy is then assessed against an independent reference sample.

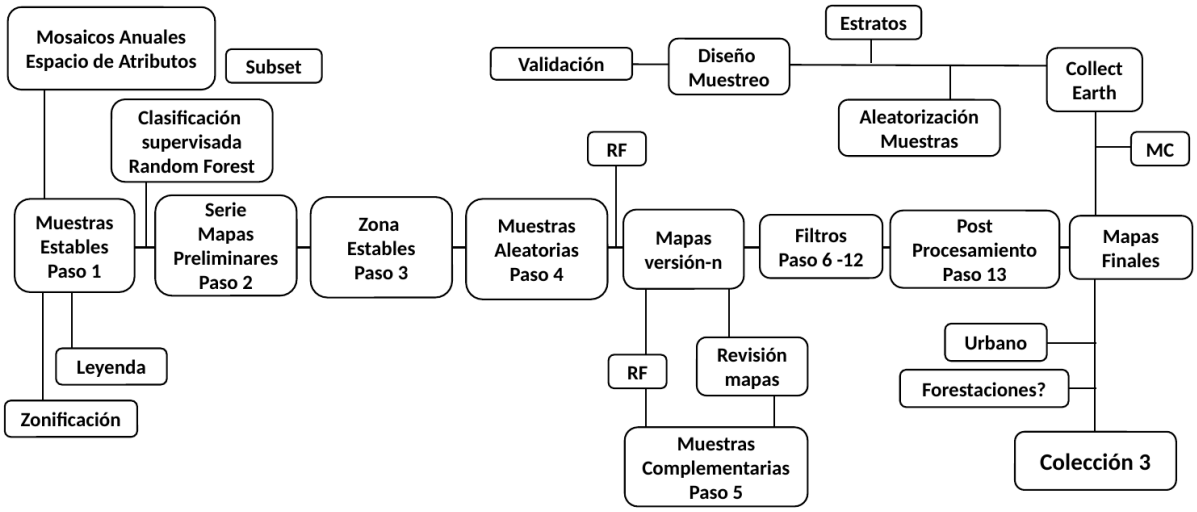

Figure 2. Methodological scheme of MapBiomas Patagonia Collection 3. Steps 1 and 2 (legend, zonification, manually digitized stable samples and preliminary map series) were inherited from Collection 2; production in Collection 3 starts at Step 3.

3. Remote sensing data

3.1. Landsat collection

The imagery used in Patagonia Collection 3 comes from the Landsat sensors Thematic Mapper (TM), Enhanced Thematic Mapper Plus (ETM+), Operational Land Imager and Thermal Infrared Sensor (OLI-TIRS) and OLI-2, on board Landsat 5, 7, 8 and 9 respectively. The Landsat collections, at 30 m spatial resolution, are accessible through Google Earth Engine and are sourced from NASA and USGS. Tier 1 surface reflectance (SR) products are used; these have undergone radiometric calibration and orthorectification based on ground control points and a digital elevation model, correcting pixel co-registration and displacement errors. For each year the best available sensor was used:

- 1986 and 1998-1999 - Landsat 5
- 2000-2002 - Landsat 7
- 2003-2011 - Landsat 5
- 2012 - Landsat 7
- 2013-2025 - Landsat 8 and 9

Between 1987 and 1997 the density and quality of Landsat acquisitions over Patagonia is insufficient to build annual cloud-free composites of comparable quality, so no classification is computed for those years. A single composite could be built for 1986, which is therefore mapped and kept as an isolated reference at the start of the series.

3.2. Landsat mosaics

Annual cloud-free composites are built from images distributed throughout the whole year. The cloud and shadow removal routine exploits the quality assessment (QA) band together with the GEE median reducer. QA values improve data integrity by indicating which pixels may be affected by artefacts or cloud contamination (USGS, 2017). In combination, the engine is instructed to select the median pixel value in the image stack, which rejects values that are too

bright (e.g. clouds) or too dark (e.g. shadows) and returns the median value of each band for each year. Mosaics are produced centrally by the MapBiomas team and consumed by the Patagonia workflow as a read-only ImageCollection.

In addition to the annual median, the mosaics carry a set of temporal reducers computed over the images of each year, which are the basis of the feature space:

- Median - median of the pixel values over the mapping period.
- Median_dry - median of the quartile of images with the lowest pixel NDVI values.
- Median_wet - median of the quartile of images with the highest pixel NDVI values.
- Amplitude (amp) - range of variation of the index over all images of the year.
- stdDev - standard deviation of all pixel values of all images of the year.
- Min / Max - lowest and highest annual value of each band.

3.3. Feature space

The full MapBiomas feature space comprises more than one hundred variables: the original Landsat bands, spectral indices, spectral-mixture fractions and textural information derived from them, each in the temporal-reducer variants listed above. Patagonia Collection 3 uses a subset of 85 of those variables (Table 1), chosen for their usefulness in discriminating the classes of the Patagonian legend. The same subset is used in all 13 zones and in all years, so that the only thing that changes between zones is the training set.

Group	Variables	n
Spectral bands (medians and variability)	blue, green, red, nir, swir1, swir2 - each as median, median_dry, median_wet and stdDev; green additionally as median_texture	25
Vegetation indices	ndvi, evi2, savi, gcvi, pri - as median, median_dry, median_wet, plus amp (ndvi, evi2) or stdDev (gcvi)	18
Water and moisture indices	ndwi (median, dry, wet, amp), wefi (median, wet, amp, stdDev), sefi (median, dry, stdDev)	11
Senescence / litter	cai - median, median_dry, stdDev	3
Spectral-mixture fractions	gv, gvs, npv, soil - median, median_dry, median_wet, amp; cloud and shade - median, max, min, amp	24
Fraction-derived index	ndfi - median, median_dry, median_wet, amp	4

Table 1. The 85-variable subset of the MapBiomas feature space used to train the Random Forest classifier in Patagonia Collection 3. Suffixes: _median = annual median; _median_dry / _median_wet = median of the driest / wettest NDVI quartile; _amp = annual amplitude; _stdDev = annual standard deviation; _texture = textural measure; _max / _min = annual extremes.

In addition to the mosaic variables, a set of ten geomorphometric descriptors is computed on the fly from the ALOS World 3D 30 m DEM (JAXA/ALOS/AW3D30_V1_1) and appended to every annual mosaic, so that they are extracted together with the spectral values at each training point (Table 2). Beyond elevation, slope and aspect, three derived variables were designed specifically for the Patagonian relief. The Heat Load Index (HLI; McCune & Keon, 2002, equation 3) summarizes the potential direct incident radiation of a site as a function of latitude, slope and folded aspect; the sign of the aspect term is inverted with respect to the original northern-hemisphere formulation, so that in the southern hemisphere north-facing slopes receive the higher load. Northing and easting are the cosine and sine of aspect weighted by a logistic function of slope, $1 / (1 + \exp(-(-5 + 0.35 \cdot \text{slope})))$, which approaches zero on flat terrain and one on slopes steeper than about 30°. This weighting solves a well-known problem of raw aspect: on nearly flat ground aspect is numerically defined but physically meaningless, and feeding it to a classifier injects noise. The weighted variables degrade smoothly to zero where aspect stops mattering.

Variable	Description
altitude	Elevation above sea level (m), ALOS AW3D30 AVE band
slope	Terrain slope (degrees)
aspect	Terrain aspect (degrees)
aspect_rad	Aspect in radians
aspect_folded	π − |aspect_rad − π|; folds aspect around the north-south axis
slope_rad	Slope in radians
latitude_rad	Absolute latitude in radians
HLI	Heat Load Index (McCune & Keon, 2002), southern-hemisphere formulation
northing	cos(aspect) weighted by a logistic function of slope
easting	sin(aspect) weighted by a logistic function of slope

Table 2. Geomorphometric variables computed on the fly and attached to the annual mosaics.

These descriptors are attached to every annual mosaic and are therefore extracted and stored together with the spectral values at each sample point. In Collection 3, however, the Random Forest is trained on the 85 spectral variables of Table 1 only, so the variables of Table 2 are not used as predictors: terrain information enters the product exclusively through the conceptual rules of Step 13 (Section 5.6). The sample tables nevertheless carry them, which makes it possible to test them as predictors in the next collection without repeating the extraction (Section 8).

4. Classification of land use and land cover

4.1. Legend

The Collection 3 legend distinguishes 14 land use and land cover classes, organized in six first-level groups (Table 3). Class identifiers and colors follow the MapBiomas national code. Not all classes occur in every zone. Three classes are not produced by the general classifier described in this document: Forest Plantations (9) is mapped with a dedicated object-based methodology, and Urban areas (24) and Glacier / Snow and ice (34) are produced by national cross-cutting working groups (Section 6).

Grp	Class (EN)	ID	Hex	Color
1.1	Forest	3	#1f8d49	
1.2	Shrubland	66	#a89358	
1.3	Flooded woody vegetation	6	#026975	
2.1	Grassland	12	#d6bc74	
2.2	Wetlands (vegas / mallines)	11	#519799	
2.3	Peatland	75	#1a88bf	
2.4	Steppe	63	#ebf8b5	
3.1	Mosaic of uses	21	#ffefc3	
3.2	Forest plantation *	9	#7a5900	
3.3	Urban area **	24	#d4271e	
4.1	High-Andean rock	29	#ffaa5f	
4.2	Deserts and badlands	25	#db4d4f	
5.1	Rivers, lakes and ocean	33	#2532e4	
5.2	Glacier / snow and ice **	34	#93dfe6	
6	Non-observed	27	#ffffff	

*Table 3. Legend of MapBiomas Patagonia Collection 3. * mapped with a dedicated methodology (Section 6.1). ** produced by a national cross-cutting working group (Sections 6.2 and 6.3).*

4.1.1. Class definitions

Forest (3). Natural woody formations dominated by trees: mainly the evergreen and deciduous Andean-Patagonian forests, and other regional forests of northern Patagonia such as the Caldén woodlands. The class is not restricted to a rigid structural definition based on height or canopy cover; it includes continuous forest as well as fragmented patches and small relicts in the forest-steppe ecotone and in drier sectors, provided they keep a spectral and ecological identity associated with natural tree cover. Tall shrub formations of valley bottoms, typically ñire (Nothofagus antarctica) stands, are also included here.

Shrubland (66). Natural woody formations dominated by shrubs or low woody plants that do not reach forest structure. Three distinct situations occur: in the forest-steppe ecotone, where it may form large continuous or fragmented areas, and inside the forest matrix where structure is low, for instance after fire; within the steppe, as small homogeneous patches taller than the surrounding shrub steppe and associated with locally moister depressions; and in the Patagonian Monte, typically north of the Río Negro and in Neuquén, as relatively continuous shrub cover. Low shrub steppe (below about one metre) is excluded.

Flooded woody vegetation (6). Natural woody formations of permanently or seasonally flooded environments, where water availability strongly conditions vegetation structure and composition. It occupies lowlands, river margins, moist depressions and river islands, and is distinguished from other wetlands by the dominance of woody rather than herbaceous or peat-forming cover. Typical examples are the islands of the Río Negro and, more rarely, lignified mallines in the steppe. It is a minor class in Patagonia and is often made up of invasive exotic trees.

Grassland (12). Cover dominated by herbaceous vegetation, especially grasses, with little or no shrub or tree participation; both natural grasslands and secondary grasslands originated by disturbance such as fire or forest removal. It occurs mainly in the north-east of the region, in some productive Andean valleys and in high-altitude environments. The class excludes grass steppe and mallines, and represents grasslands on mountain slopes, high-altitude grasslands near the vegetation limit, and naturalized grasslands of native and exotic species associated with agricultural land. Its boundaries with steppe and, in some cases, with agricultural cover are gradual.

Wetlands - vegas and mallines (11). Herbaceous or grass-dominated wetlands on saturated soils or soils with high water availability, in both steppe and forest environments. They are distinguished by high productivity, hygrophilous vegetation and contrast with the drier surrounding matrix, and usually show a spectral signal of high moisture and seasonal or permanent greenness. They occur on relatively flat, depressed landscapes where water accumulates, may contain open water bodies, and do not normally occur on slopes. They are distinguished from peatlands especially in Tierra del Fuego.

Peatland (75). Wetlands with accumulation of organic matter and peat development, typical of the cold, humid south of Patagonia, especially Tierra del Fuego and southern Santa Cruz. Their ecological dynamics differ from vegas and mallines: prolonged water saturation, specialised vegetation, and conditions that favour accumulation of soil organic carbon.

Steppe (63). The steppe formations typical of extra-Andean Patagonia: low, open and discontinuous vegetation with variable combinations of grasses, low shrubs and bare soil. It is a very extensive and heterogeneous class that integrates different structural and functional states within a single large regional ecological unit. Its boundaries are diffuse both with shrubland,

where woody development increases, and with non-vegetated areas, where plant cover drops sharply; it is separated from deserts and badlands by having more than 5 % vegetation cover.

Mosaic of uses (21). Landscapes transformed by agricultural and livestock activities in which different uses and covers coexist at small spatial scales. In Patagonia it is associated mainly with smallholdings, productive Andean valleys and irrigated areas of the steppe, combining crops, pastures, shelterbelts, fallow plots, rural infrastructure and other anthropogenic covers. Because of this heterogeneity these situations are not separated into more specific classes at this stage of the mapping.

Forest plantation (9). Areas planted with trees for production, mainly commercial plantations of exotic species, especially pines. It is characterised by relatively regular spatial arrangements and by growth trajectories typical of forest management, and is harder to identify in early stages, when canopy cover is low and the spectral signature resembles other covers. It includes wildling invasions spread from the plantations.

Urban area (24). Surfaces occupied by urban infrastructure, settlements, built-up areas and other artificial covers, with a high density of roads and non-vegetated surfaces. Vegetated (green) areas contained within the city are considered urban, whereas equivalent areas on the urban fringe are not.

High-Andean rock (29). High-mountain surfaces with very low or no vegetation cover - although with a significant presence of high-Andean flora, including endemics - dominated by exposed rock, outcrops and scree. Although it shares the criterion of sparse vegetation with other non-vegetated areas, it is distinguished by its topographic, climatic and biogeographic context.

Deserts and badlands (25). Non-vegetated surfaces, or surfaces with extremely sparse vegetation (below 5 % cover), in arid and semi-arid environments: badlands, deserts, dunes, deflation basins, salt flats, alluvial plains and other landforms dominated by bare soil, loose sediments or surface crusts. It may grade into the more open steppe.

Rivers, lakes and ocean (33). Permanent surface water bodies: lakes, ponds, rivers and the ocean. Identification relies on the low near-infrared reflectance characteristic of water. In training, lakes and ponds - as large as possible - were prioritized in order to avoid edge effects and the spatial variability of river courses.

Glacier / snow and ice (34). Glaciers, ice and persistent snow, mainly in the Andes and the southern Patagonian ice fields. The class represents permanent or very long-lasting cryospheric cover, as distinct from transient or seasonal snowfall.

Non-observed (27). Pixels for which land cover cannot be reliably determined because of limitations in satellite observation, in Patagonia mainly topographic shadows in the Andes. These areas do not correspond to a land cover type but to a condition of insufficient observation. Training samples are generally not collected for this class, but it is used in post-processing to remap classes wrongly identified under deficient spectral signal, for instance water detected on shaded slopes.

4.2. Classification algorithm and parameters

Classification is performed zone by zone and year by year with the Random Forest algorithm (Breiman, 2001) implemented in Google Earth Engine (ee.Classifier.smileRandomForest). The same hyper-parameters are used in every zone and year: 100 trees, 4 variables tried per split, a minimum leaf population of 25 pixels, and a fixed random seed (1) so that runs are reproducible.

Requiring at least 25 samples in a terminal node limits the depth of the trees and avoids leaves grown to purity, which stabilizes the class-probability output used in Step 13.

A single classifier is trained per zone and year: the training table of that year carries the class label of the sample point and the spectral values of that year's mosaic at that point, and the trained model is applied to the mosaic of the same year. Because the sample points are fixed across the series and only the spectral values change, the classifier sees a consistent definition of each class throughout the whole period, which is what makes the resulting time series comparable.

4.3. Stable areas (Step 3)

Collection 3 does not repeat the manual digitization of temporally invariant polygons that opened Collections 1 and 2. Instead, the stable-area mask is derived directly from the final, fully post-processed Collection 2 maps. This has two consequences worth stating explicitly: the stable areas already incorporate the criteria embedded in the Collection 2 post-processing, and the interpreters' effort in Collection 3 could be concentrated on reviewing maps and collecting complementary samples rather than on building the initial training set from scratch.

For each of the 13 zones, the 28 annual layers of the Collection 2 final classification (1986 and 1998-2024) are stacked and two statistics are computed per pixel: the number of distinct classes over the series, and the modal class. A pixel is considered stable when it satisfies two conditions simultaneously: it took fewer than five distinct classes over the 28 years, and its modal class was assigned in at least 24 of those years (about 85 % of the series). Stable pixels are labelled with their modal class; all other pixels are masked. The result is exported per zone as PAT-STABLE-REGION-X.

The same script computes how many training samples each class should receive in each zone. The relative area w of each class within the zone is computed on the reference year (the last year of the Collection 2 series), and the number of samples is allocated as n = round(200 + 6000·w). This gives a floor of 200 samples for classes that are rare in the zone and a ceiling of about 2,000 samples for a class occupying 30 % or more of the zone, so that abundant classes are well represented without swamping the rare ones. The allocation table is exported alongside the stable map and consumed by the next step.

4.4. Random stable samples (Step 4a)

Sample points are drawn by stratified random sampling over the stable map of each zone, using the class-wise allocation computed in Step 3, at 30 m resolution and with a fixed seed. The same point set is then used for all years: for each of the 29 mapped years (1986 and 1998-2025) the annual mosaic - with the geomorphometric bands attached - is sampled at every point, and the resulting table (class label plus feature values) is exported as samples-stable-<year>-<zone>. Points that fall on a no-data pixel of the mosaic are discarded. Keeping the geometry of the sample fixed while letting the spectral values vary with the year is what makes the per-year classifiers mutually consistent.

4.5. Stable pre-classification (Step 4b)

A first Random Forest classification is produced for every zone and year using only the stable samples. The 29 annual results of a zone are stacked into a single multi-band image, one band per year, and exported as PAT-REGION-X. This product - referred to as the "stable pre-classification", or [stable] layer - is the object the interpreters review in Step 5. It deliberately

contains no complementary samples and no post-processing, so that the errors that remain are attributable to the training set rather than to the filters.

4.6. Complementary samples and final classification (Step 5)

The stable pre-classification is reviewed zone by zone and hexagon by hexagon in the GEE Code Editor. Each reviewer has, for every year, three map layers: [stable], the pre-classification under review; [Col2], the published Collection 2 map with all its post-processing, which shows how persistent problems were resolved by filters in the previous collection; and [NEW], the classification recomputed on the fly as complementary samples are added. The false-color annual mosaics and the high-resolution imagery of the GEE basemap are available as backdrops, and clicking on a pixel plots its NDVI time series together with its [stable] and [NEW] class trajectories. Comparing [stable] with [NEW] shows the effect of the samples just added; comparing them with [Col2] shows what the filters can be expected to fix on their own.

Complementary samples are digitized as polygons, drawn on the Landsat imagery even when high-resolution imagery is used for interpretation, and stored in a geometry layer named after the class (c03, c66, …) carrying a "class" property. They must meet the same criteria as the original stable samples: they must be temporally stable over the whole period, internally homogeneous even when the class itself is heterogeneous, relatively small (of the order of 5 × 5 pixels), and distributed so as to cover the spatial range of the class within the cell. For natural vegetation classes the last years of the series are checked with particular care, to make sure no transformation occurred; for anthropogenic classes the first years are checked instead. Three questions guide the decision: which class would I want the map to show in this micro-zone; is this micro-zone a good training example for that class; and does this sample add internal heterogeneity that earlier samples do not already cover, without drifting onto the spectral boundary of the class.

The complementary samples digitized for Collection 2 are loaded into the script as a reference. They were not used to building the Collection 3 pre-classification, so they act as a record of where a previous interpreter judged the map to be wrong. Where such a sample agrees with the [stable] map it is ignored; where it disagrees and the current reviewer considers it a good correction, it is promoted into the Collection 3 complementary set.

Random points are then generated inside the digitized polygons by stratified random sampling, with a ceiling of 1,000 points per class and zone. That ceiling was lowered in the two zones where the available polygons were scarce: in Zone 4 to 200 points for flooded woody vegetation and wetlands and to 500 for shrubland and grassland, and in Zone 12 to 500 for wetlands. These points are merged with the stable samples of the corresponding year and the Random Forest is retrained and re-run for all 29 years. The cycle of review, resampling and reclassification is repeated, each iteration producing a new version of the per-zone classification (PAT-ClassComplement-REGION-X-vN), until the reviewers judge that the remaining errors cannot be corrected with additional samples. Two cautions are applied throughout: some classification problems are structural and cannot be solved with samples, and an excessive number of samples for one class produces unwanted changes elsewhere in the map.

Three classes are removed from the stable sample set at this stage and are therefore not predicted by the general classifier: Forest plantations (9), Glacier / snow and ice (34) and Non-observed (27). The first two are produced by dedicated workflows (Section 6) and the third is reintroduced, where appropriate, by the post-processing rules of Step 13. Urban areas (24) are likewise never sampled. Forest plantations are nevertheless reviewed during sample collection,

but only to correct commission errors: where forest has been classified as plantation, complementary forest samples are added to remove those pixels, accepting the resulting omission because plantations are mapped separately.

4.7. Class probabilities and the second most probable class

In parallel with the hard classification, a second Random Forest with identical parameters is trained in multi-probability mode. For each pixel and year it returns the vector of per-class votes; the script pairs each probability with its class identifier, sorts the pairs by decreasing probability and keeps the top three. The result is stored as an array band per year holding a 3 × 2 array - probability and class identifier for the first, second and third most probable classes - and exported per zone as PAT-ClassCompMultiProb-REGION-X. By construction the first row reproduces the hard classification.

These products are integrated across zones in the same way as the hard classification (Section 5.1), producing a Patagonia-wide layer of the second most probable class per pixel and year. Their purpose is to support the conceptual post-processing of Step 13: when a rule establishes that the class assigned to a pixel is implausible in its geomorphic context, the second most probable class provides an evidence-based replacement instead of a single class fixed a priori by the rule. Three of the Step 13 rules are built this way - water and wetlands on steep, high terrain, and mosaic of uses at altitude (Section 5.6.1) - and they are the rules that apply over the largest areas, which is why an informed replacement matters. It should be kept in mind that Random Forest vote shares are informative but not calibrated posterior probabilities; they are used here as a ranking, not as a probability estimate.

5. Post-processing

Because the classification is pixel-based and performed independently for each year, the raw output contains two kinds of artefact: spatial noise (isolated pixels and very small patches that do not correspond to real cover units) and temporal noise (implausible one-, two- or three-year excursions of a pixel to a different class and back). Post-processing removes both, then applies a set of conceptual rules that encode expert knowledge the classifier cannot express. The chain is applied in a fixed order, each step consuming the output of the previous one (Table 4).

Step	Filter	Purpose
6	Integration and gap fill	Merge the 13 zones; fill no-data from the temporally nearest valid class
7	Spatial filter	Mode filter on patches of ≤ 6 connected pixels
8	Temporal filter, 3-year window	Correct one-year anomalies
9	Temporal filter, 4-year window	Correct two-year anomalies
10	Temporal filter, 5-year window	Correct three-year anomalies
11	First and last year	Correct 1998 and 2025 against their neighbours
12	Final spatial filter	Remove artefacts created by the temporal filters; reinsert 1986
13 a	Conceptual reclassification	Probability-, elevation- and sector-based rewrite rules
7-12 bis	Second filtering pass	The spatial and temporal filters of Steps 7-12 re-applied to the reclassified series
13 b	Final assembly	Perennial water bodies; forest plantations; extension to 1985-1997; residual code check

Table 4. Post-classification chain of Collection 3.

5.1. Integration and gap fill (Step 6)

The 13 per-zone classifications are merged into a single Patagonia-wide image, one band per year. Two integrations are built. In the first, each zone classification is clipped to its own polygon expanded by 30 m and the overlapping stack is reduced with a per-pixel minimum, so that where two adjacent zones disagree in their narrow overlap the lower class code wins - a deterministic and reproducible tie-break rather than an arbitrary order of precedence. In the second, the same operation is performed on the zonification buffered by 5 km. The wide integration is then blended underneath the narrow one, so that it only shows through where the narrow one has no data, which fills any sliver gaps left between adjacent zone polygons.

A gap-fill filter is then applied to the no-data values that remain - pixels for which no valid classification exists in a given year, typically because the mosaic had no valid observation there. No-data values are replaced by the temporally nearest valid classification: the series is first traversed forward, each year inheriting from the previous one where it is empty, and then backwards, so that gaps at the start of the series are filled from the future. A gap therefore survives only if the pixel was no-data throughout the entire series. Finally, connectedPixelCount is computed for every year and appended as an additional band per year, to be consumed by the spatial filter. The result is exported as PAT-INTEGRACION_GapFill.

The same script exists in a multi-probability variant, which integrates the first and second most probable class layers of Section 4.7 across the 13 zones and applies the same gap-fill and connectivity computation to them.

5.2. Spatial filter (Step 7)

The spatial filter is designed to remove isolated pixels and very small patches without eroding the edges of genuine cover units. It combines two operations. First, connectedPixelCount identifies, for each pixel and year, how many contiguous neighbours share its class. Second, a mode filter is computed over a 5 × 5 pixel square window. The mode value is written back only where the connected patch has six pixels or fewer, and the original classification is preserved everywhere else. Patches larger than six pixels are therefore never modified, and the minimum mapping unit implied by the filter is of the order of 0.5 ha. The result is exported as PAT-INTEGRACION-Spatial_Filter.

5.3. Temporal filters (Steps 8-10)

The temporal filters detect transitions that are implausible over short time spans and replace them with the class the pixel held before and after the excursion. They operate on the continuous series 1998-2025; the isolated 1986 map is excluded, because it has no temporal neighbours, and is reinserted in Step 12. Three window sizes are applied in sequence, each consuming the output of the previous one.

Three-year window (Step 8). For each class c and each central year t between 1999 and 2024, a pixel is flagged when its class at t equals c while its classes at t − 1 and t + 1 are both different from c. The flagged pixel at t is then replaced by its class at t − 1. This removes one-year excursions to and from a class.

Four-year window (Step 9). The same logic is extended to two-year excursions: a pixel is flagged when its classes at t and t − 1 both equal c while its classes at t + 1 and t − 2 differ from c, and both flagged years are replaced by the class at t − 2. Because a single pass over a fixed set of anchor years cannot cover every possible position of a two-year anomaly, the filter is applied

twice, once with even-numbered anchor years and once with odd-numbered ones, the second pass taking the output of the first as its input.

Five-year window (Step 10). Three-year excursions are corrected in the same way, with three interleaved groups of anchor years chained one after the other, so that every possible phase of a three-year anomaly is covered.

In all three filters the correction is applied only when the replacement value belongs to a list of admissible classes. In Steps 8 and 9 that list is [3, 66, 6, 12, 11, 75, 63, 21, 29, 25]; in Step 10 it is the shorter list [3, 12, 11, 63, 21], so the longest anomalies are only corrected when they are bounded by the most extensive classes. Water (33) is deliberately absent from every list, so that the filters never convert land into water; Forest plantations (9), Non-observed (27) and Glacier / snow and ice (34) are absent because they are not produced by the general classifier. Within each filter the classes are processed in the fixed order 3, 63, 66, 12, 21, 6, 11, 75, 29, 25, which gives priority to the more extensive and better-defined classes; the result of each class pass feeds the next.

5.4. First and last year (Step 11)

The first and last years of the series have only one temporal neighbour and are therefore not protected by the moving-window filters. Two dedicated rules are applied. For 1998, if a pixel differs from class c while both 1999 and 2000 equal c, the 1998 value is set to c; this prevents spurious regeneration or loss in the first year of the series. For 2025, if a pixel equals c in both 2023 and 2024 but differs in 2025, the 2025 value is set to c. The classes are processed in the order 3, 66, 12, 11, 75, 6, 63, 21, 29, 25 for the first year. For the last year, Forest (3) is deliberately excluded from the rule, so that a genuine loss of forest in the final year - typically caused by fire - is preserved rather than filtered away.

5.5. Final spatial filter and reinsertion of 1986 (Step 12)

The temporal filters can create new small patches where a corrected pixel is surrounded by uncorrected ones. A final spatial filter with the same parameters as Step 7 - mode over a 5 × 5 window, applied only where the connected patch has six pixels or fewer - is therefore applied at the end of the chain. At this point the 1986 band, which was held out of the temporal filters, is reinserted from the output of Step 7 and included in the final spatial filtering, so that the 1986 map receives the same spatial treatment as the rest of the series. The result is exported as PAT-INTEGRACION-Filtros.

5.6. Conceptual reclassification (Step 13)

The statistical filters described above act on the geometry and the temporal behaviour of the classification, but they cannot express ecological or geomorphic knowledge about where a class can and cannot occur. Step 13 encodes that knowledge as an explicit, auditable set of reclassification rules. Every rule has the same form: within a defined domain - a combination of slope and elevation thresholds, a zone, or a hand-digitized sector - pixels of a given class are rewritten as another class. The rules are applied as a chain, each one taking the output of the previous one as its input, so that the order in which they are written is the order in which they take effect and the state of the map after any rule can be inspected. They are applied to every year of the series simultaneously, which keeps the correction temporally consistent. Terrain descriptors - elevation, slope and aspect - are computed from the NASADEM 30 m digital elevation model.

The rules fall into three families, described in the next three subsections. The reclassified series is exported as PAT-INTEGRACION-FiltConcept and then passed through the filter chain a second time (Section 5.7) before the final assembly (Section 5.8).

5.6.1. Rules resolved with the second most probable class

Three of the rules identify a context in which the assigned class is geomorphically implausible but in which no single alternative class is obviously correct. In those cases the pixel is rewritten with the second most probable class returned by the classifier for that pixel and year (Section 4.7), which lets the replacement follow the local evidence instead of a class fixed a priori.

- Water (33) on slopes steeper than 5° and above 1,000 m of elevation. Standing water cannot occupy steep, high terrain; the signal is almost always a topographic shadow or a snow patch.
- Wetlands - vegas and mallines (11) - on slopes steeper than 5° and above 1,000 m. Vegas and mallines occupy flat, depressed landscapes where water accumulates, so a moist, high-productivity signal on a steep high slope is something else, most often a shrub or grass formation.
- Mosaic of uses (21) above 1,000 m in Zones 2 and 4. Agricultural and peri-urban mosaics do not occur at those elevations in those sectors.

These three rules replaced the fixed rewrites used in Collections 1 and 2, where water on slopes became Non-observed and wetlands and mosaic of uses on slopes became Shrubland regardless of context. Adding the elevation condition also restricts the rules to the Andean sector, where the confusion actually occurs, instead of applying them over the whole territory.

5.6.2. Elevation rules: High-Andean rock and Deserts and badlands

High-Andean rock (29) and Deserts and badlands (25) are spectrally almost identical - both are essentially bare surfaces - and are separated by their geomorphic context rather than by their reflectance. The separation is therefore imposed by an elevation threshold, which varies along the latitudinal gradient because the altitudinal belts descend southwards. The rule is applied in both directions, so that the elevation threshold defines the boundary between the two classes exactly.

Sector	Elevation threshold
Zones 1-9 (north)	1,150 m
Zones 10-12 (south)	800 m
Zone 13 (Tierra del Fuego)	300 m

Table 5. Elevation thresholds separating High-Andean rock (29) from Deserts and badlands (25). Below the threshold, pixels classified as 29 are rewritten as 25; at or above it, pixels classified as 25 are rewritten as 29.

5.6.3. Sector rules

A set of hand-digitised polygons delimits sectors in which a specific, persistent confusion was identified by the interpreters and could not be resolved with complementary samples. Within each polygon - in some cases further restricted to one zone - a single class-to-class rewrite is applied. Most of these sectors were inherited from the review of Collections 1 and 2; many of the corrections that those collections needed are no longer required, either because the class in question is now mapped correctly or because it is resolved by one of the rules above or by a cross-cutting layer. Table 6 lists the rules that remain in Collection 3, in the order in which they are applied.

Rule	Rewrite	Domain	Rationale
Wetlands → Mosaic of uses	11 → 21	Zone 9	Irrigated valleys around Sarmiento, where the wetland signal corresponds to agricultural use
Mosaic of uses → Steppe	21 → 63	Atlantic coast	Coastal steppe systematically read as agricultural mosaic
Wetlands → Shrubland	11 → 66	Atlantic coast	Coastal shrubland read as vega
Flooded woody vegetation → Shrubland	6 → 66	Atlantic coast	Coastal shrubland read as flooded woody vegetation
Forest → Shrubland	3 → 66	Zone 1	Low and open woody formations over-predicted as forest
Forest → Wetlands	3 → 11	Zone 1	Estuarine mallines read as forest
Grassland → Shrubland	12 → 66	Zone boundary	Shrubland read as grassland along a zone seam
Steppe → Grassland	63 → 12	Zone 4	Grassland read as steppe at a zone seam
Steppe → Grassland	63 → 12	Zone 5	Grassland read as steppe at a zone seam
Grassland → Shrubland	12 → 66	Zone 5	Shrubland read as grassland
Shrubland → Steppe	66 → 63	Zone 8	Steppe read as shrubland
Flooded woody vegetation → Forest	6 → 3	Zone 4	Persistent commission of flooded woody vegetation over forest

Table 6. Sector-based reclassification rules of Step 13, in order of application. "Domain" is the zone or sector to which the hand-digitised polygon belongs.

5.7. Second filtering pass (Steps 7-12 re-applied)

The conceptual rules of Section 5.6 rewrite whole sectors and whole classes at once, and in doing so they can introduce exactly the kind of artefact that the post-processing chain is designed to remove: small residual patches along the edge of a rewritten sector, and pixels whose trajectory becomes inconsistent because the replacement value differs between neighbouring years. The reclassified series is therefore passed through the whole filter chain a second time - gap fill and connectivity, spatial filter, the three temporal filters, the first- and last-year correction, and the final spatial filter - with the same parameters as the first pass. The result is exported as PAT-INTEGRACION_FCo-Filtros and is the input to the final assembly.

5.8. Final assembly (Step 13, second part)

Three operations complete the product.

Perennial water bodies. The temporal filters are deliberately forbidden to create water, which protects the map from spurious lakes but leaves genuine permanent water bodies flickering between Water and other classes in years with poor observations. Within the polygons of the national perennial continental water layer larger than the minimum size retained (complemented with a small set of hand-digitised bodies), any pixel classified as Water (33) in five or more years of the series is therefore set to Water in every year. The rule is deliberately confined to mapped perennial bodies, so it cannot invent water anywhere else.

Temporal extension to 1985-1997 and forest plantations. The 1986 map is replicated over 1985 and 1987-1997, so that the published collection presents a continuous annual series of 41 years from 1985 to 2025. These replicated years carry no independent observation and should be read as a constant extrapolation of the 1986 state, not as evidence about land cover dynamics in that period. The forest plantation layer (Section 6.1) is then written into the series as class 9 wherever it is present, superseding the general classification for that class.

Residual code check. As a last safeguard, every pixel and year is checked against the 14 codes of the legend. Any value outside that set - which can survive from the integration of the zones or from the chained rewrites - is replaced by the modal valid class of that pixel over 1998-2025,

or by Non-observed (27) if the pixel has no valid class in the whole series. The corrected series is exported as PAT-FINAL and is the product distributed as MapBiomas Patagonia Collection 3.

6. Cross-cutting and complementary mappings

Three classes of the legend are not produced by the general classifier described in Sections 4 and 5. They are mapped by dedicated workflows, with their own methodologies and their own input data, and are integrated into the product afterwards. For that reason their samples are excluded from the training set of the Patagonia classifier and the corresponding class codes are excluded from the replacement lists of the temporal filters, so that the general workflow neither predicts them nor overwrites them. Forest plantations are integrated within the Patagonia workflow itself (Section 5.8); urban areas and glaciers are integrated at the national assembly stage, when the regional products of all MapBiomas Argentina regions are merged into the national collection.

6.1. Forest plantations

An assessment carried out for Collection 3 against the National Forest Plantation Inventory showed that the general pixel-based approach substantially under-classifies rainfed forest plantations: over stands established in 2000-2002, the area strictly classified as Forest plantation remained below 7 % during the first decade after planting and reached at most 19 % by 2022, with most of the remaining area assigned to Shrubland, Steppe and, increasingly with stand age, Forest. Two mechanisms explain this: during establishment and juvenile growth the canopy is too open for a 30 m pixel to be dominated by conifer foliage, and once the canopy closes the spectral separation between exotic conifer plantations and native forest is weak.

Collection 3 therefore maps rainfed plantations with a dedicated Geographic Object-Based Image Analysis (GEOBIA) workflow in Google Earth Engine. Its main features are a seasonal feature space - spectral bands and indices computed as summer, autumn and spring medians, plus phenological differences between seasons, designed to exploit the contrast between deciduous native forest and evergreen plantations - GLCM texture metrics that capture the geometric regularity of plantation rows, topographic variables derived from FABDEM, segmentation into superpixels with the SNIC algorithm, and a Random Forest classification of the resulting objects. The full description is given in ANNEX I, which should be read as part of this ATBD.

The plantation layer is written into the map at the final assembly of Step 13, after the whole post-processing chain has run (Section 5.8), so that the filters can neither erode it nor propagate it in time. During complementary sample collection, interpreters review this class only to remove commission errors, accepting the resulting omissions because the dedicated layer supersedes the general classifier for this class.

6.2. Urban areas

Urban areas (24) are mapped by a national cross-cutting working group with a methodology common to all Argentine regions and to all years of the collection, and are written into the maps when the regional products are assembled into the national collection. The corresponding code therefore does not appear in the intermediate products of the Patagonia workflow. No training samples are collected for this class in Patagonia, although the class is included in the accuracy assessment. Where an urban area is not covered by the cross-cutting layer, the general classifier

assigns it to Deserts and badlands, Steppe, Mosaic of uses, Shrubland or Forest, depending on the surrounding vegetation context.

6.3. Glaciers, snow and ice

Glacier / snow and ice (34) is likewise addressed by a cross-cutting national initiative, which maps the class for the whole country with a dedicated classification of its own, and its samples were removed from the Collection 3 training set. As with urban areas, the layer is written into the maps at the national assembly stage rather than within the Patagonia workflow, so the code does not appear in the intermediate Patagonia products. The National Glacier Inventory and the cross-cutting glacier layer are used as auxiliary information during sample collection and map review, to keep interpreters from adding samples of other classes over glaciated surfaces.

7. Accuracy assessment

Unlike Collections 1 and 2, Collection 3 incorporates a probabilistic accuracy assessment based on an independent reference sample. The sample is allocated by stratified random sampling, using the classification of a mid-series year (2006) as the stratification layer. Sample sizes are computed from the standard expressions for a target standard error of overall accuracy, balanced by the mapped area of each class and then truncated, so that extensive classes do not absorb the whole sample and rare classes still receive enough points to support a class-wise estimate; a fixed increment is added for minority classes. Point locations are randomized with a minimum-distance restriction to limit spatial autocorrelation between reference units. Five reference years are evaluated: 1986, 1999, 2006, 2015 and 2023.

Reference labels are assigned by visual interpretation in Collect Earth, using Landsat time series, high-resolution imagery and spectral index trajectories. Each point is interpreted independently by two interpreters, and a third interpreter resolves disagreements. From the resulting confusion matrices, overall accuracy, user's and producer's accuracy per class, and area estimates with confidence intervals are computed following good-practice recommendations for area estimation from stratified samples. The design, the estimators and the results are described in full in the MapBiomas Argentina validation ATBD, to which the reader is referred

9. References

Achanta, R., & Süsstrunk, S. (2017). Superpixels and polygons using simple non-iterative clustering. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 4651-4660.

Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.

Burkart, R., Bárbaro, N., Sánchez, R. O., & Gómez, D. A. (1999). Ecorregiones de la Argentina. APN, PRODIA. 43 pp.

Haralick, R. M., Shanmugam, K., & Dinstein, I. (1973). Textural features for image classification. IEEE Transactions on Systems, Man and Cybernetics, SMC-3(6), 610-621.

Hawker, L., Uhe, P., Paulo, L., Sosa, J., Savage, J., Sampson, C., & Neal, J. (2022). A 30 m global map of elevation with forests and buildings removed. Environmental Research Letters, 17(2), 024016.

Lencinas, J. D., Salvaré, F., Barberá, I., Dieguez, G., Masera, P., Bagnato, C., Moyano, J., Mateo, C., Gowda, J., Verón, S., Banchero, S., Leszczuk, A., & Eljall, A. (2026). Methodology for Classifying Rainfed Plantations in Patagonia. MapBiomas Argentina, Collection 3.

McCune, B., & Keon, D. (2002). Equations for potential annual direct incident radiation and heat load. Journal of Vegetation Science, 13(4), 603-606.

Olson, D. M., Dinerstein, E., Wikramanayake, E. D., Burgess, N. D., Powell, G. V. N., Underwood, E. C., … Kassem, K. R. (2001). Terrestrial ecoregions of the world: a new map of life on Earth. BioScience, 51(11), 933-938.

Oyarzábal, M., Clavijo, J., Oakley, L., Biganzoli, F., Tognetti, P., Barberis, I., Maturo, H. M., Aragón, R., Campanello, P. I., Prado, D., Oesterheld, M., & León, R. J. C. (2018). Unidades de vegetación de la Argentina. Ecología Austral, 28(1), 40-63.

Roy, D. P., Kovalskyy, V., Zhang, H. K., Vermote, E. F., Yan, L., Kumar, S. S., & Egorov, A. (2016). Characterization of Landsat-7 to Landsat-8 reflective wavelength and normalized difference vegetation index continuity. Remote Sensing of Environment, 185, 57-70.

USGS (2017). Landsat Quality Assessment Band. United States Geological Survey.

ANNEX I: Methodology for Classifying Rainfed Plantations in Patagonia

Lencinas José Daniel, Salvaré Fernando, Barberá Iván, Dieguez Gonzalo, Masera Pablo, Camilo Bagnato, Jaime Moyano, Camila Mateo, Juan Gowda, Verón Santiago, Banchero Santiago, Leszczuk Andres, Eljall Ana.

1. Introduction

1.1. Scope and content of the section

This section provides an overview of the satellite image processing steps, input feature layers, and classification workflows used to map rainfed forest plantations in Patagonia as part of the Collection 3 Land Cover Classification. It describes the underlying theory, justifications, and procedures used to generate the baseline map for 1986 and subsequent annual forest plantation maps for the 1998–2025 period.

1.2. Frame and accuracy analysis of the Coleccion 1 and 2

In Patagonia, the total area of rainfed plantations established with species of the genera *Pinus* and *Pseudotsuga* (*Pseudotsuga menziesii* (Mirb.) Franco) is estimated at 106,679 ha and 1,430 ha, respectively (SAGYP, 2026). Within the genus *Pinus* (*Pinus ponderosa* Dougl. ex Laws., *P. contorta* subsp. *murrayana* (Grev. & Balf.) Critchf., *P. radiata* D. Don, and *P. jeffreyi*), *P. ponderosa* accounts for 81.8% of the total planted area (INPF, 2017). The afforestation rate peaked at 5,000 ha/year during the 1990s, later declining to 1,000 ha/year until 2015 (Bava et al., 2015). Over the last decade, the planting rate in Patagonia has roughly hovered around an average of 500 ha per year.

It should be noted that national government subsidies provided during the 1970s and 1980s supported exotic conifer tree planting in Andean Patagonia within native forest stands (Bava et al., 2015). By the mid-1990s, promotion policies shifted toward areas free of native forests, though degraded native shrublands were still targeted for conifer planting. From 2000 onward, forest plantations have been established almost exclusively in non-forested sites.

Unlike Collection 3 (1986, 1998–2025), which incorporates a dedicated validation framework, Collection 1 (1986, 1998–2023) and Collection 2 (1986, 1998–2024) lacked independent validation datasets. Consequently, a quantitative evaluation using independent reference data was conducted to support the interpretation of the classification accuracy for Collections 2. To this end, polygon-based samples were selected from the National Forest Plantation Inventory (NFPI) database (INPF, 2017), focusing on rainfed forest plantations established in 2000 and 2002, which feature the most reliable plantation age attributes. This high reliability is due to the administrative process for receiving subsidies, which requires applicants to submit site and area records that are subsequently verified through field monitoring by the provincial and national authorities.

For validation, a dataset of 10,134 plantation stands, represented as polygons, was spatially overlaid with the MapBiomas Collection 3 Stable Pre-classification product, derived of the Collection 2, to quantify the area intersecting the Forest Plantations (ID 9), Forest (ID 3), Shrublands (ID 66), Grassland (ID 12), and Steppe (ID 63) classes. Figure 1 shows the mean

annual percentage of the reference plantation-stand area classified into each of these classes from the year of plantation establishment through 2023.

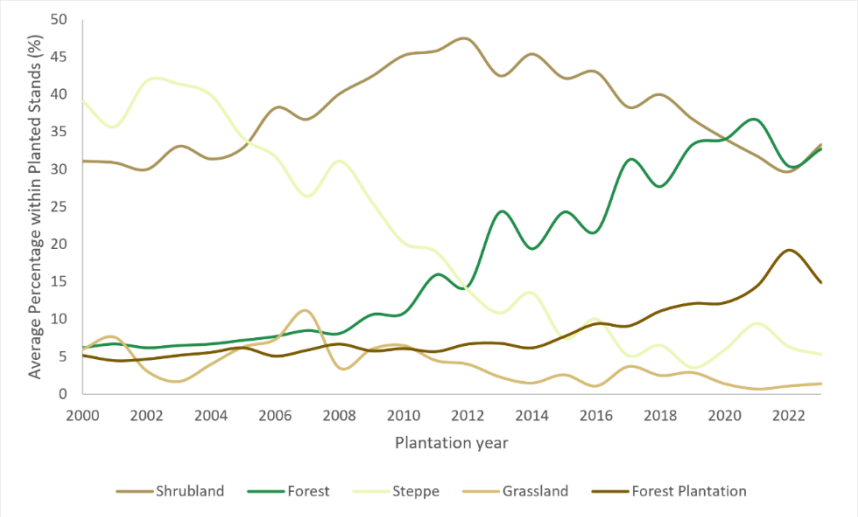

Fig. 1. Evolution of the average percentage of MapBiomas land-cover classes within planted forest stands established between 2000 and 2002.

During the first decade following planting, the surface area strictly classified as Forest Plantation (ID 9) remained below 7%, whereas the combined Shrubland (ID 66) and Steppe (ID 63) classes accounted for nearly 70%. This pattern is directly linked to the establishment and juvenile growth phases of the stand: during the first ten years, conifer foliage canopy cover is insufficient for full closure, causing the spectral signature of a 30 m Landsat pixel to be dominated by bare ground, steppe, or adjacent shrub cover.

From 2015 onward—as the forest stands reached 13 to 15 years of age—the proportion classified as Forest Plantation increased continuously, rising from 7.6% in 2015 to a peak of 19.1% in 2022. This trend indicates that the MapBiomas classification methodology requires stands to reach an intermediate developmental stage with higher canopy closure before reliably discriminating them as rainfed forest plantations. Concurrently, the relative cover of the Steppe and Grassland classes declined. Pixels that failed to be classified as plantation were predominantly assigned to Shrubland and Forest.

Notably, Forest (ID 3) expanded steadily from 6.2% in 2000 to 32.7% by 2023. By the end of the time series in 2023, the combined detection of Forest Plantation (14.9%) and Forest (32.7%) reached 47.5%. This demonstrates that while the algorithm effectively detects dense tree cover once stands mature into their second decade, its capacity to spectrally differentiate exotic forest plantations from native forests remains limited.

Several methodological limitations should be noted. The analysis was restricted to stands aged 20–23 years (representing ~35% of the total planted area), while stands older than 30 years were excluded due to unreliable or missing age data. Additionally, all database polygons were assumed to be fully stocked. In practice, the reported boundaries often include non-stocked areas such as internal roads, unmapped gaps, or localized establishment failures where

replanting was unsuccessful. Despite these limitations, the evaluated sample clearly demonstrates that the Forest Plantation class is significantly under-classified. Consequently, these findings prompted the development of a novel methodological approach to classify Patagonian rainfed forest plantations.

2. Methodology

In this section, we present a tailored methodology for mapping rainfed forest plantations across the Patagonian region. As part of Collection 3, we implemented a Geographic Object-Based Image Analysis (GEOBIA) workflow within Google Earth Engine (GEE). Unlike traditional pixel-based approaches, this methodology adopts spatial objects (superpixels) as the primary units of classification. Objects are initially delineated from a selected feature space using the Simple Non-Iterative Clustering (SNIC) segmentation algorithm (Achanta and Süsstrunk, 2017) and subsequently classified via a Random Forest model across a seven-phase framework (Fig. 2).

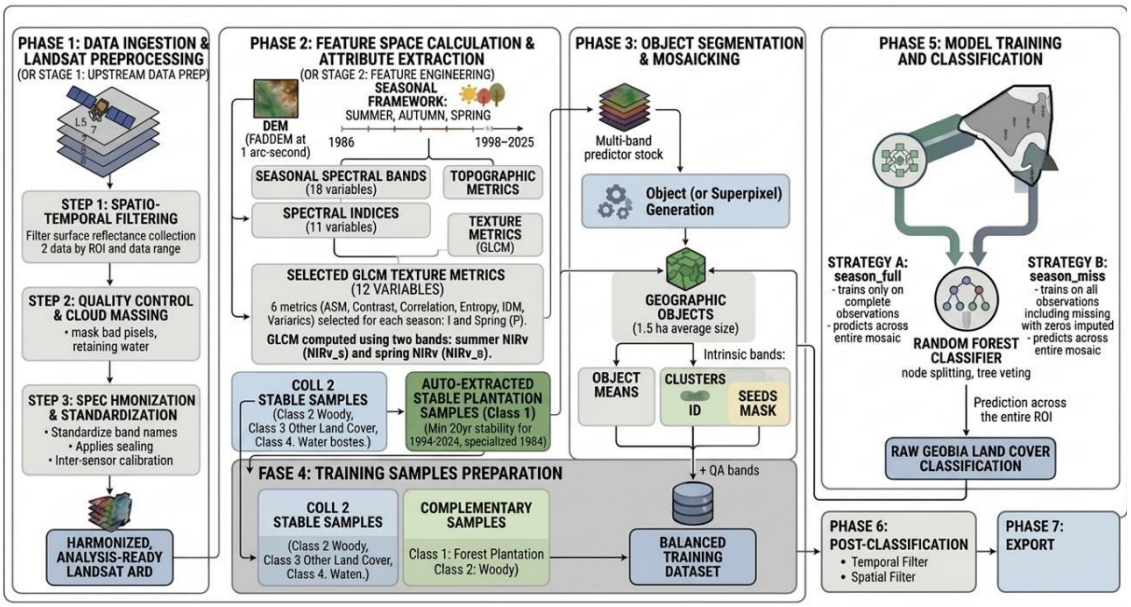

Fig 2. Patagonian rainfed forest plantation mapping workflow. Figure prepared by the authors with assistance from OpenAI's ChatGPT (GPT-5.6).

2.1. Data Ingestion and Landsat Image Preprocessing

Step 1. Landsat Spatio-Temporal Filtering and Data Ingestion: Surface Reflectance Collection 2 data from Landsat 5 (TM), Landsat 7 (ETM+), Landsat 8 (OLI), and Landsat 9 (OLI-2) were filtered by the region of interest (ROI) and the target temporal range. The ROI encompasses the entire area with potential for rainfed tree plantations, spanning native forests, forest–steppe ecotones, and steppe ecosystems. The study area is geographically bounded by 71.9095° W, 36.4152° S (upper-left corner) and 70.2699° W, 45.2460° S (lower-right corner),

covering approximately 145,431 km² and extending 148 km in the east–west direction and 982 km in the north–south direction.

Step 2. Quality Control and Cloud Masking: Bad or corrupted pixels were masked out using the QA_PIXEL quality band. Specifically, dilated clouds, cirrus, high-confidence clouds, cloud shadows, and snow/ice were excluded. Water pixels (bit 7) were intentionally retained at this step to enable downstream multi-class land cover classification.

Step 3. Spectral Harmonization and Standardization: Band names across all sensors were standardized to a unified naming scheme (blue, green, red, nir, swir1, and swir2). Radiometric scaling parameters were applied, and inter-sensor spectral cross-calibration was performed following Roy et al. (2016).

2.2. Feature Space Calculation and Attribute Extraction

Built upon a seasonal framework, the methodology evaluates a comprehensive feature space comprising primary spectral bands, derived spectral and textural indices, and topographic metrics. Unlike the general classification approach used in previous MapBiomas Patagonia Land Cover Collections (Collections 1 and 2), feature calculation for forests plantations classification was structured across three key seasons—summer, autumn, and spring—for each year of the multi-temporal time series (1986 and 1998–2025). This seasonal strategy was designed specifically to enhance the spectral discrimination between deciduous native forests (predominant across Patagonian ecosystems) and evergreen forest plantations.

The complete feature space consists of:

- Seasonal Spectral Bands (18 variables): Six core Landsat spectral bands (blue, green, red, nir, swir1, swir2) computed as seasonal medians across summer, autumn, and spring.
- Spectral and Textural Indices (11 variables): Target vegetation and canopy condition metrics computed across the seasonal framework.
- Topographic Metrics: Terrain descriptors derived from the FABDEM (Forest And Buildings removed Copernicus DEM) digital elevation model. FABDEM was selected as the bare-earth Digital Terrain Model (DTM) at 1 arc-second (~30 m) spatial resolution (projects/sat-io/open-datasets/FABDEM). Hawker et al. (2022) generated this product by applying Random Forest models trained on ICESat-2 ATLAS and GEDI spaceborne LiDAR to remove canopy and building height biases from the Copernicus GLO-30 DEM. Recent evaluations in forested mountain catchments indicate that FABDEM significantly outperforms SRTM, MERIT DEM, and Copernicus GLO-30, representing the current state-of-the-art elevation dataset for dense forest cover (Marsh et al., 2023).

All spectral bands and derived indices were computed as seasonal median composites at a uniform 30 m spatial resolution.

Table 1. Feature Space Composition: Seasonal spectral bands, indices, and topographic metrics

Band / Index	**Landsat**		**Temporal Seasons**		
	5 (TM) & 7 (ETM+)	**8 (OLI) & 9 (OLI-2)**	**Summer 01.01/04.01**	**Autum 04.01/07.01**	**Spring 08.15/12.22**
Blue	B1	B2	✓	✓	✓
Green	B2	B3	✓	✓	✓
Red	B3	B4	✓	✓	✓
NIR	B4	B5	✓	✓	✓
SWIR 1	B5	B6	✓	✓	✓
SWIR 2	B7	B7	✓	✓	✓
NIRv (Near-Infrared Reflectance of Vegetation)	NDVI * NIR		NIRv-v	-	NIRv-p
EVI2 (Enhanced Vegetation Index 2)	2.5* $(NIR-Red)/(NIR+2.4Red+1)$		EVI2-v	*EVI2-o**	EVI2-p
NDMI (Normalized Difference Moisture Index)	$(NIR-SWIR1)/(NIR+SWIR1)$		NDMI-v	*NDMI-o**	NDMI-p
Delta EVI2 (summer-spring)	(EVI2-v – EVI2-p)	dEVI2vp			
Delta EVI2 (summer-autumn)	(EVI2-v – EVI2-o)	dEVI2vo			
Delta NDMI (summer-spring)	(NDMI-v – NDMI-p)	dNDMIvp			
Delta EVI2 (summer-autumn)	(NDMI-v – NDMI-o)	dNDMIvo			
FABDEM (30 m)	Altitud				
FABDEM (30 m)	Slope (grad)				

- ✓ calculated, - no calculated, * indices served exclusively as input data for calculating the autumn seasonal delta.

Texture metrics

To capture the spatial heterogeneity and geometric regularity characteristic of forest plantations, texture metrics derived from the Gray-Level Co-occurrence Matrix (GLCM, Haralick et al., 1973) were incorporated into the feature space composition. Plantation geometry, row alignment and spacing, generates structural patterns that contrast with native forest cover or forest-shrubland-steppe transition sites. Unlike mature native forests or those with slower development phases, spectral texture evolves dynamically with plantation age and silvicultural interventions (e.g., thinning and pruning), altering canopy discontinuity and projected shadows in the spectral data.

To generate the texture variables, the GLCM algorithm was applied by selecting two synthetic bands, NIRv (nirv_v and nirv_p). Prior to this, an 8-bit scaling was performed on both bands. GLCM was configured to run with the Texture(size: 2) parameter, which defines the neighborhood radius (in number of pixels) around the central pixel to build the co-occurrence matrix. That is, the texture was calculated within a 5×5 pixel moving window centered on the evaluated pixel, determining local variation patterns.

The GLCM algorithm automatically computes 18 texture metrics, six of which were selected for both the spring and summer seasons: Angular Second Moment, Inverse Difference Moment, Contrast, Correlation, Entropy, and Variance. Figure 3 illustrates the selected texture metrics alongside a conceptual framework for their interpretation. As higher spatial resolution datasets are integrated, the relevance of these metrics is expected to grow significantly.

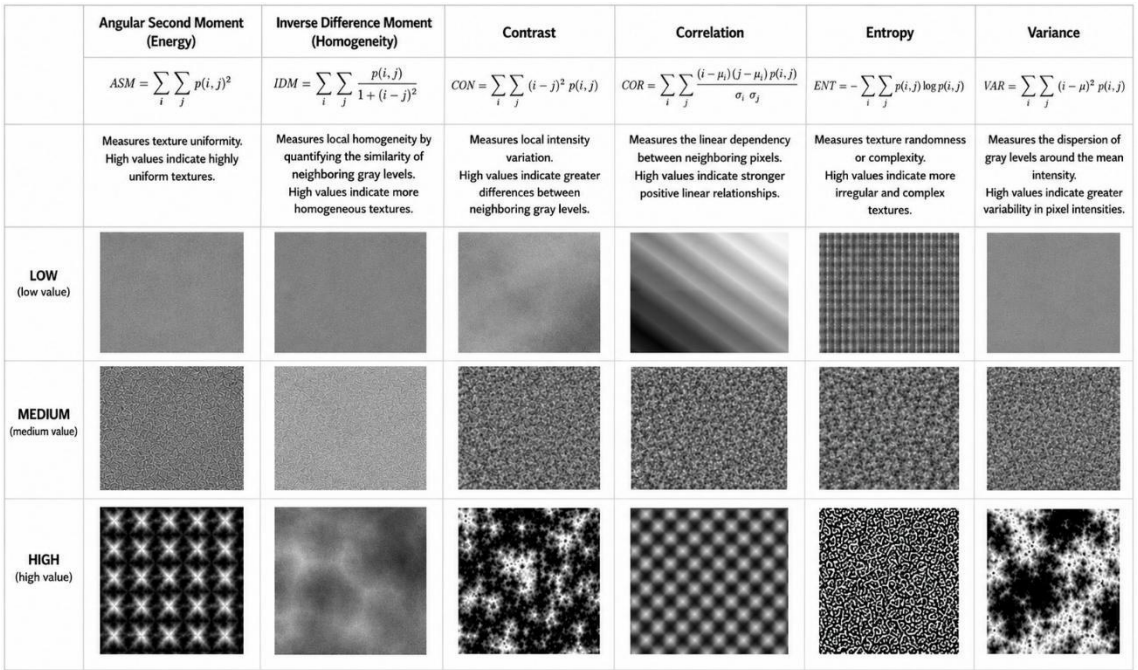

Fig. 3. Selected texture feature space attributes. Created by the authors with assistance from OpenAI's ChatGPT (GPT-5.6).

2.3. Object Segmentation and Mosaicking

Geographic Object-Based Image Analysis (GEOBIA) was performed using the Simple Non-Iterative Clustering (SNIC) algorithm (Achanta and Susstrunk, 2017) implemented in Google Earth Engine (ee.Algorithms.Image.Segmentation.SNIC). SNIC is a superpixel segmentation algorithm that groups spatially contiguous pixels with similar spectral, textural, and topographic characteristics into homogeneous objects. Unlike traditional iterative clustering methods, SNIC uses a priority queue to grow clusters around seed points in a single pass. This design significantly reduces computational overhead and memory consumption while eliminating local overlap issues across large spatial domains.

The SNIC was ran over the 42-band composite containing spectral reflectances, indices, phenological deltas, GLCM textures, and topographic variables (Table 2). The SNIC parameters were set as follows: size: 5, compactness: 0.8, connectivity: 4, and neighborhoodSize: 50. Setting size to 5 places initial seed points at 5-pixel intervals, which corresponds to a 150 m grid spacing for 30 m Landsat data. Pilot tests across diverse plantation sites demonstrated that this configuration yields an average object size of approximately 1.5 ha over forests planted. A compactness of 0.8 emphasizes geometric regularity, driving segment boundaries toward circular or compact polygonal shapes. The 4-neighbor connectivity restricts pixel aggregation strictly to shared edges (top, bottom, left, and right), fostering the creation of spatially cohesive clusters. Lastly, setting neighborhoodSize to 50 bounds the search distance relative to the seed points, effectively mitigating boundary artifacts during tile or patch processing.

Table 2. Spectral, index, texture and topographical bands integrated for object segmentation.

Block / Season	Input Spectral Variables and Indices for SNIC	Band Count
Summer (_v)	blue_v, green_v, red_v, nir_v, swir1_v, swir2_v, nirv_v, evi2_v, ndmi_v	9
Spring (_p)	blue_p, green_p, red_p, nir_p, swir1_p, swir2_p, nirv_p, evi2_p, ndmi_p	9
Autumn (_o)	blue_o, green_o, red_o, nir_o, swir1_o, swir2_o	6
Phenological Deltas	delta_evi2_vp, delta_ndmi_vp, delta_evi2_vo, delta_ndmi_vo	4
GLCM Texture	Summer (nirv_v): nirv_v_asm, nirv_v_contrast, nirv_v_corr, nirv_v_ent, nirv_v_idm, nirv_v_var Spring (nirv_p): nirv_p_asm, nirv_p_contrast, nirv_p_corr, nirv_p_ent, nirv_p_idm, nirv_p_var	12
Topography	elevation, slope	2

SNIC calculates the mean value of every input variable across all pixels belonging to each objects. As after segmentation pixels carried the mean value of its cluster in each band,

classification is computed pixel-wise, but all pixels in a cluster get the same class. Besides SNIC bands we also stored the number of valid observations per pixel and season, with values varying by pixel, even if they belong to the same cluster. The valid observations count was used in the classification phase.

The segmentation and mosaicking step resulted in a multi-band image per year (1986 and 1998-2025) that was used for extracting training samples (phase 4) and for classifying the whole region of interest (phase 5).

2.4. Training data

Due to strict project timelines, the methodology prioritized the synthesis and integration of existing datasets of stable samples over primary data collection.

The optimized approach for establishing the training samples consisted of a hybrid framework aimed at ensuring high thematic confidence. On one hand, stable samples were selected from Collection 2 for pixels corresponding to non-forest plantation classes. On the other hand, specific training samples for the 'Forest Plantation' class were automatically extracted using an independent reference dataset for support.

From the Collection 2 dataset, stable samples were extracted across the 1985–2024 time. A pixel was defined as "stable" if it maintained the same remapped land-cover class across a minimum historical threshold ($frequency \geq 20$ years). This procedure was applied to select stable samples for the following aggregated categories:

Class 2. Woody: Forests and shrublands (excluding forest plantations)

Class 3. Other land cover: Herbaceous vegetation, pastures, agriculture, and

urban areas

Class 4. Water bodies: Rivers, lakes, and ice/snow

Through this workflow, a total of 5,000 stable samples were selected for each of these three classes, distributed across the Region of Interest (ROI). Note, however, that each stable sample is used as training data in all years.

Stable samples for the 'Forest Plantation' class (Class 1) were derived by first filtering the National Plantation Forest Inventory database to include only polygons with reliable stand age data. Stable pixels were subsequently extracted within these benchmark boundaries from the Collection 2 time series (1998–2024). Pixels were classified as 'stable' if they maintained the target class for at least 20 years, evaluated via an attribute table recording annual binary values (1 = Forest Plantation, 0 = non-plantation).

Specifically for the year 1986, NPI polygons were selected if, following database curation and visual interpretation, they could be inferred as forest plantations with high probability, restricting selection to polygons representing stands at least 10 years old by 1986. Stable

plantation samples were subsequently generated manually within these polygons using the 1986 imagery mosaic as visual support. It should be noted that thematic accuracy for this target year is lower, as plantations established during the 1970s and 1980s were predominantly embedded within native forest matrices—many of which lacked silvicultural management—thus making spectral discrimination particularly challenging.

2.4.1. Complementary samples

A preliminary Random Forest classification was performed to assess the need for additional training data; subsequent visual inspection against Landsat and high-resolution imagery in GEE guided the collection of complementary samples.

Utilizing the NFPI, updated provincial forest plantations inventory for Río Negro province, and high-spatial-resolution imagery, additional training samples were selected for both the 'Forest Plantation' and 'Woody classes. A tailored protocol was developed to carry out the collection of complementary samples. This strategy resolved class confusion and broadened representation, as the initial stable plantation samples were predominantly located outside native forest areas. These complementary samples for 'Woody vegetation' and 'Forest Plantation' were applied across the entire time series and restricted to the 2010–2025 period, respectively.

To mitigate class imbalance between the 'Forest Plantation' and 'Woody' classes, a balancing strategy was implemented to restrict their sample size difference to a maximum of 10%, downsampling whichever class was more abundant.

1998–2024: The training dataset comprised both stable and complementary

samples, with the latter restricted to the Forest Plantation class from

2010 onward. Samples were subsequently filtered to retain only those meeting the criteria of having been classified as 'Forest Plantation' in Collection 2 for the corresponding year. Combining the stable and complementary samples, 6273 samples were used for class 1, 5585 for class 2, 5000 for class 3, and 5000 for class 4.

2025: Classification relied on the stable training samples derived from 2024.

1986: 78 Training samples were drawn exclusively from the 1986 dataset.

2.5. Model trainnig and Classification

A Random Forest classifier was trained for each year, using the corresponding samples. However, we used different sets of samples among years, due to data missingness. Season-wise medians could not be computed in areas and periods with less than one valid observation. This issue occurs frequently across large regions prior to the year 2000 due to sparse Landsat coverage, as well as in recent years over the Andean Mountain range, where autumns are characterized by high cloud cover and snow. Missing pixels are assigned a value of zero by the SNIC algorithm. Although Random Forest (RF) can ingest these zero-filled inputs, an abundance of zeros reduces the information gain of the affected seasonal variables during

tree node splitting. Nevertheless, if the remaining available seasons carry sufficient spectral signal, the RF model can still achieve strong performance.

Following an iterative visual quality assessment, a dual-model training and prediction strategy was established using two distinct dataset configurations:

1. **season_full**: the Random Forest is trained using only the samples with at least one valid observation in each season (zero-filled values are removed). Prediction is then executed unconditionally across the entire mosaic, even in zero-filled pixels.
2. **season_miss**: the Random Forest is trained using all samples in the year, retaining the zero-filled pixels. Prediction is likewise performed over the entire seasonal mosaic.

Predictions for the 1998–2000 period used the season_miss dataset, while season_full was selected for 1986 and from 2001 onward. We also tried a different training/prediction scheme to deal with missing data: we built a mosaic of annual-level summaries to avoid missing data at the season-level. However, that model performed worse than the approaches described above, based on visual evaluation of predictions. This suggests that ignoring the seasonal patterns is worse than including the zero-filled missing data for training and/or prediction. No water masking was applied during prediction because the training dataset included robust, stable samples for the water class. However, this class was not used to map surface water in the final product; only the extracted plantation layer is incorporated into the final LULC map.

The Random Forest for each year was trained using the following hyperparameters: 100 trees, bag fraction = 0.5 (bagFraction = 0.5), and the standard defaults were set for the variables per split ($\sqrt{M}$), minimum leaf population = 1, and maximum number of nodes uncapped.

2.6. Post-Classification

Visual inspection of the classification results revealed that, overall, the plantation classification exhibited very low omission errors. However, there was a noticeable overestimation (commission error), which was more pronounced in the western region, specifically in plantation sites located within native forest areas. Consequently, temporal and spatial filters were designed and iteratively calibrated to mitigate commission errors (Table 3).

Table 3. Temporal and spatial filters.

Filter 1 **Single-Year Spike Removal Filter** **1998 to 2025**	For each year k, the filter evaluates whether the pixel is classified as Forest Plantation (Class 1). If both the preceding year (k-1) and the subsequent year (k+1) are classified as non-plantation (≠1), this isolation condition is triggered, and the label for year k is reassigned to the Forest class.
Filter 2 **Temporal persistence filter** **1998 to 2025**	The algorithm calculates the maximum length of consecutive years in which a pixel was classified as *Forest Plantation* ($N_{consec} \geq 5$). Pixels that fail to meet this temporal continuity threshold are rejected, ensuring that only established plantations are retained for subsequent filtering steps.
Filter 3 **Epoch-Based Commission Correction Filter**	To correct persistent commission errors in native forest transition zones, this filter evaluates historical class dynamics across two distinct epochs: pre-2010 and post-2010. If a pixel was dominantly classified as Forest prior to 2010, a strict post-2010 persistence threshold is enforced ($P \geq 12$ years). Pixels falling below this threshold are classified as false positives and reverted to Forest.
Filter 4 **Minimum Mapping Unit (MMU)**	It implements a threshold to remove spatial noise (the "salt-and-pepper" effect) by filtering out small, isolated pixel clusters and retaining only contiguous plantation patches that meet a minimum area of 0.5 hectares

Due to an 11-year data gap between 1986 and 1998, the 1986 baseline classification was excluded from temporal filtering (Filters 1–3). Instead, it was processed exclusively through a spatial Minimum Mapping Unit (MMU) filter of 0.5 ha (5 connected pixels) to remove isolated patches.

After applying filters 1, 2, 3, and 4, a binary multi-band ee.Image (uint8: 1 = plantation, 0 = non-plantation) clipped to the ROI is generated. Finally, a dynamic user interface was implemented to integrate pixel-level inspection tools, synchronize historical fire and auxiliary reference layers, and generate temporal presence charts (1985–2025) upon clicking any pixel (Fig. 4).

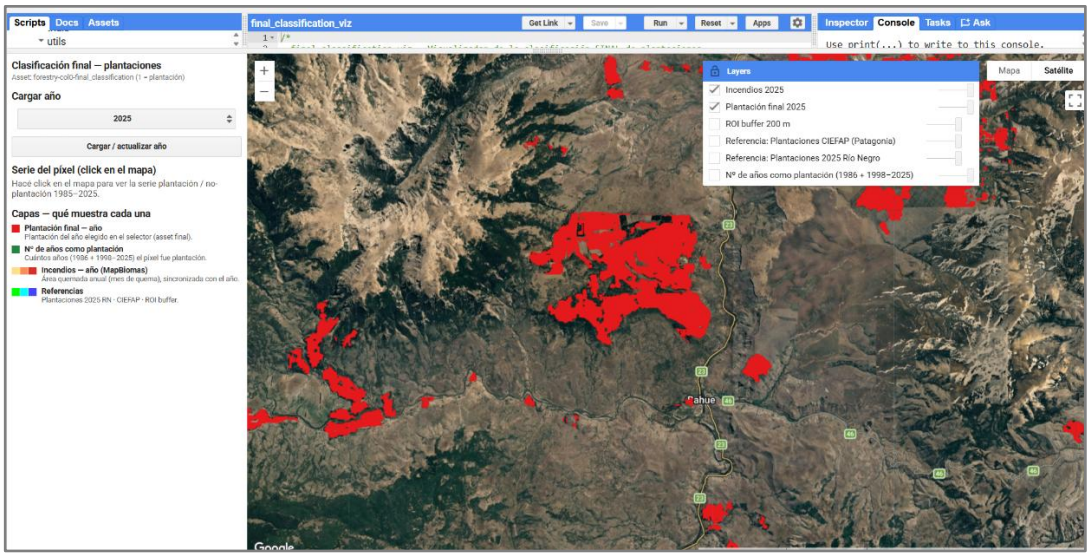

Fig. 4. Interactive user interface for visualization and pixel-level time-series inspection of forest

plantations (1985–2025)

2.7. Export

In order to incorporate the plantation classification into the Collection 3 LULC product, the plantation class was isolated and remapped from Class 1 to Class 9. Spatial export was executed across a 260,240 ha final ROI encompassing all plantation sites in Patagonia. This domain was defined by taking the classified plantation extent, applying a 200 m buffer, and refining the boundary using auxiliary data from forest plantation inventories updated with newly detected plantation areas (Fig. 5).

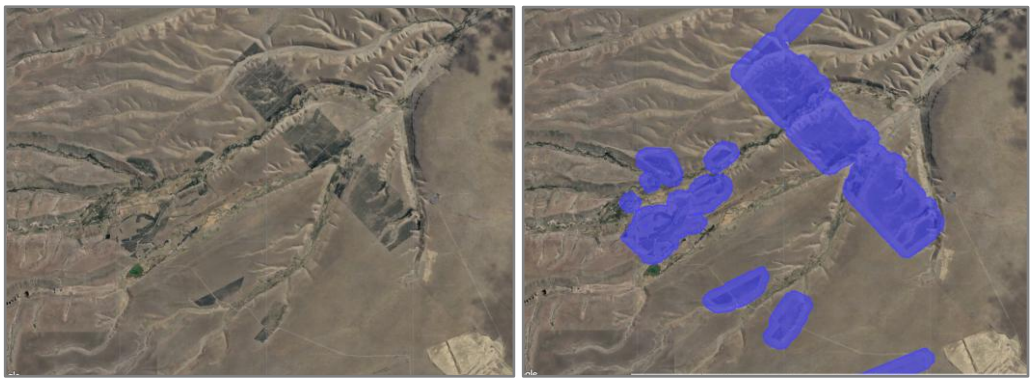

Fig. 5. Final ROI for exporting the classified forest plantation class. Left: Forest plantation stands.

Right: Final ROI boundary.

References

Achanta, R., & Susstrunk, S. (2017). Superpixels and polygons using simple non-iterative clustering. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 4651–4660). IEEE. https://doi.org/10.1109/CVPR.2017.495

Bava, J., Loguercio, G., & Salvador, G. (2015). ¿Por qué plantar en Patagonia? Estado actual y el rol futuro de los bosques plantados. Ecología Austral, 25(2), 101–111. https://doi.org/10.25260/EA.16.25.2.0.155

Haralick, R. M., Shanmugam, K., & Dinstein, I. (1973). Textural features for image classification. IEEE Transactions on Systems, Man, and Cybernetics, SMC-3(6), 610–621. https://doi.org/10.1109/TSMC.1973.4309314

Hawker, L., Uhe, P., Paulo, L., Sosa, J., Savage, J., Sampson, C., & Bates, P. (2022). A 30 m global map of elevation with forests and buildings removed. Environmental Research Letters, 17(2), Article 024016. https://doi.org/10.1088/1748-9326/ac4d4f

Marsh, C. B., Harder, P., & Pomeroy, J. W. (2023). Validation of FABDEM, a global bare-earth elevation model, against UAV-lidar derived elevation in a complex forested mountain catchment. Environmental Research Communications, 5(3), Article 031009. https://doi.org/10.1088/2515-7620/acc56d

INPF (2017). Inventario de plantaciones forestales en secano región Patagonia. Ministerio de Agroindustria, UCAR; CIEFAP.https://www.magyp.gob.ar/sitio/areas/desarrollo-foresto-industrial/foresto-ndustria/_archivos4/000000_Inventario%20Forestal/000000_Inventario%20Patagonia%20-%20Andina.pdf

Roy, D. P., Kovalskyy, V., Zhang, H. K., Vermote, E. F., Yan, L., Kumar, S. S., & Egorov, A. (2016). Characterization of Landsat-7 to Landsat-8 reflective wavelength and normalized difference vegetation index continuity. Remote Sensing of Environment, 185, 57–70. https://doi.org/10.1016/j.rse.2015.12.024

SAGYP (2026). Tablero de plantaciones forestales. Secretaría de Agricultura, Ganadería y Pesca. Ministerio de Economía de la Nación. Recuperado el 23 de agosto de 2026, de https://www.magyp.gob.ar/desarrollo-foresto-industrial/tablero-plantaciones-forestales.php