



Algorithm Theoretical Basis Document (ATBD)

Appendix - Glaciers

MapBiomas Water – Ecuador Collection 4, Version 1

EcoCiencia Executive Director
Carmen Josse

MapBiomas Ecuador Coordinator
Wagner Holguín

MapBiomas Water Ecuador Coordinator
Juan Espinosa

September 2026

Table of contents

1. Introduction	3
1.1. Scope and Use of the Data.....	4
2. Study area	4
3. Imagery mosaics.....	5
4. Classification.....	6
4.1. Spectral Mixture Analysis (SMA)	8
4.2. Spectral Classification Logic (Threshold Hierarchy).....	8
4.3. Classification Variables	9
5. Post-Classification	11
5.1. Gap Filling	11
5.2. Spatial filter	12
5.3. Persistence Filter.....	14
5.4. Consistency filter	15
5.5. Corrective filter	16
6. Discussion and Data Limitations	17
7. Conclusions	19
8. References	20

1. Introduction

The present Algorithm Theoretical Basis Document (ATBD) comprises an adaptation of the methodology proposed by the MapBiomass Water and Glaciers team of Peru.

The document provides a detailed description of the processing chain applied to satellite images corresponding to wet-season conditions or to periods immediately following high-precipitation events (the dataset referred to as “wet”), within the group of snow and glacier cover detection algorithms for the Ecuadorian territory. The purpose of using images with these “wet” characteristics is to generate an annual time series of maps representing the maximum extent of snow cover for each year during the 1985–2025 period, using as the main input surface reflectance mosaics derived from the Landsat constellation (TM, ETM+, and OLI).

It is important to note that this product should not be interpreted as a map of permanent glaciers. The mapped snow cover includes both exposed glacial ice and fresh or seasonal snow deposited over the Andean summits during periods of increased solid precipitation. In the tropical latitudes of the Ecuadorian Andes, such seasonal snow is a meteorologically sporadic and rapidly melting phenomenon; therefore, its detection through optical sensors requires a specific methodological approach aimed at maximizing the probability of capturing these events during peak accumulation.

The general strategy of the snow and glacier mapping method is based on a dual seasonal-criteria approach:

Wet criterion (“wet”): Documented in this ATBD, it detects the maximum extent of snow-covered surfaces by leveraging reflectance percentiles associated with fresh snow, wet snow, and saturated ice. It includes both glaciers and seasonal snow over rock outcrops along the flanks of volcanic edifices.

Dry criterion (“dry”): Processed in parallel and complementarily, it uses reflectance percentiles corresponding to dry-season conditions, when seasonal snow has melted and only the actual glacial ice mass remains exposed.

The spatial and temporal comparison between the wet and dry products ultimately enables the discrimination of seasonal snow cover from the permanent glacier mass, forming the basis for generating glacier maps for the Andean region.

Therefore, the algorithms described in the following sections (spectral classification, spectral mixture analysis, temporal gap filling, spatial filters, persistence filters, consistency filters, and corrective filters) should be understood as components of the wet criterion of the system, whose purpose is to ensure that the maximum snow signal is robustly captured while minimizing noise and information gaps without compromising the spatial and temporal coherence of the series.




Figure 1. Methodological flowchart for the generation of snow-cover maps. Source: MapBiomass Peru.

1.1. Scope and Use of the Data

Given that snow cover in tropical latitudes exhibits high seasonal variability and a rapid response to climatic events, the use of individual satellite images is severely limited by persistent cloudiness. To mitigate this constraint, the algorithm operates on annual composite mosaics calculated from surface reflectance medians. This strategy allows synthesizing the dominant spectral signal of the landscape, reducing residual atmospheric noise and contamination from pixels affected by undetected clouds or shadows.

Snow-estimation data provide a reference for the horizontal extent and altitudinal ranges reached by solid precipitation over glacier-bearing volcanoes. This information is essential for distinguishing accumulated snowfall episodes from the glacier mass itself.

The datasets produced by MapBiomass Agua, within its glacier and snow modules, present limitations in detecting “black ice” or debris-covered glaciers, due to the spectral decomposition complexity associated with pixels in these areas.

2. Study area

The study area (Figure 2) was defined based on the Randolph Glacier Inventory (RGI Consortium, 2017, v.0), corresponding to the entire Andean region of Ecuador. Within this region, glacier zones were delineated using an altitudinal buffer of 500 meters below the elevation of the glaciers detected in 1985, in order to establish the classification area.




Figure 2: Study area

3. Imagery mosaics

Snow classification relied on Landsat image mosaics specifically generated for mapping the maximum extent of solid precipitation events. These mosaics were constructed from pixels belonging to wet-season images or images meeting high-moisture criteria following precipitation events. To build each mosaic, a percentile-based distribution of pixel values from the filtered images within the analysis region was used.

The pixel analysis follows a percentile distribution, where pixels corresponding to the 75th percentile are selected for composing the wet-characteristic (“wet”) mosaics.

The wet-characteristic mosaic captures conditions representing an estimated maximum annual snow cover. It is constructed by selecting images where the NDSI value is greater than or equal to the 75th percentile of the annual distribution. From these filtered images, the pixel-wise median is computed, generating the bands referred to as “wet.” The 75th percentile was chosen to balance the maximization of captured snow cover with the minimization of noise introduced by clouds typical of the wet season.

The percentile behavior in image selection operates pixel by pixel over the temporal distribution of NDSI values. For a pixel to be included in the 75th-percentile mosaic as snow cover, it must have been classified as snow in approximately 25% or more of the valid images of the year, corresponding to the statistical position of the 75th percentile in the ordered distribution of values.




Figure 3: Construction of wet-season mosaics based on percentile distribution.
Source: MapBiomias Peru.

4. Classification

The algorithms in this stage of the workflow aim to generate annual binary maps of snow and exposed ice presence/absence for the Ecuadorian territory, using the Landsat satellite time series (TM, ETM+, and OLI) as the primary input.

The snow-classification process operates on the wet-characteristic mosaic (75th percentile of the annual NDSI series), whose purpose is to capture the maximum snow cover of the year. The classification of the Landsat mosaics was performed entirely on the Google Earth Engine platform using threshold-based rules. Snow-cover classification from the wet-characteristic mosaic does not employ a supervised classifier such as Random Forest. Instead, it relies exclusively on spectral threshold rules applied directly to the mosaic variables through conditional where operators in the programming language. This classification approach responds to the temporal instability of seasonal snow, which prevents the construction of a sufficiently stable training-pixel set for the wet season.

A set of sequential conditions is applied to the mosaics, where all pixels begin with an initial assignment of “no snow cover,” and each threshold rule may overwrite the result of the previous one. The first condition evaluates the reflectance in the near-infrared band of the wet-characteristic mosaic, discarding as snow cover those pixels that do not exceed the minimum threshold established for each sensor.

Similarly, the second condition evaluates reflectance in the red band of the wet-characteristic mosaic, applying differentiated thresholds for Landsat 5, Landsat 7, Landsat 8, and Landsat 9, again discarding pixels that do not meet these criteria.

The final identification of snow cover is based on a dual condition that must be met simultaneously:

- The annual minimum NDSI value, representing the lowest snow-cover signal recorded for that pixel throughout the year, must be greater than or equal to 100 on the scaled index (equivalent to a real NDSI ≥ 0.0), and
- The snow fraction derived from Spectral Mixture Analysis (SMA) must be greater than or equal to 20 percent.

Pixels satisfying both conditions are assigned to the snow-presence class.




Figure 4: Empirical decision tree combined with Random Forest for snow classification. Source: MapBiomass Peru.

Finally, as a quality-control step, pixels with residual cloud contamination are identified using the cloud fraction derived from SMA. Those exceeding the threshold value are classified as “invalid data” and removed from the final classification, ensuring that only spectrally reliable observations are included in the snow-cover product.

In summary, snow and ice exhibit extremely high reflectance in the visible spectrum (especially blue and red) and very low reflectance in the shortwave infrared (SWIR). However, the exact spectral response varies according to the band configuration and radiometric calibration of each Landsat mission. The median near-infrared parameter acts as a preliminary radiometric filter. If a pixel shows reflectance in the near-infrared (NIR) or red band below these thresholds, it is automatically discarded as snow (or labeled as “non-snow”). This is based on the fact that snow, to be considered as such in a median mosaic, must maintain high reflectance; low values suggest the presence of water, vigorous vegetation, or dark rock.

Additionally, a Digital Elevation Model (SRTM) and its derivative, slope, are integrated. The inclusion of slope is critical for snow dynamics because on very steep slopes (>15–20 degrees), reflectance can be affected by illumination geometry (slope-orientation effects).

Furthermore, snow tends not to accumulate stably on extremely steep slopes or in deeply shaded valleys, unless it corresponds to debris-covered glaciers or ice within glacial cirques.

4.1. Spectral Mixture Analysis (SMA)

After preprocessing, the workflow calls the Spectral Mixture Analysis (SMA) module. This module generates the fractions *gv* (green vegetation), *npv* (non-photosynthetic vegetation), *soil*, *cloud*, and *snow*.

In the Tropical Andes, a 30 × 30-meter Landsat pixel is rarely “pure snow.” It is typically a mixture of ice, exposed rock (debris), debris-covered glaciers, and scattered snow patches. Spectral mixture analysis allows modeling the pixel as a linear combination of pure spectral signatures (endmembers). The snow fraction (0–100%) indicates the proportion of the pixel covered by snow, providing a far more robust metric for glacier-retreat studies than a simple hard classification.

4.2. Spectral Classification Logic (Threshold Hierarchy)

The final classification is constructed through a cascade of logical conditions (where statements).

Initially, all pixels are assumed to be “Non-Snow.” They remain classified as “Non-Snow” if their NIR reflectance is below the threshold. Although snow in the near-infrared (NIR) is not as bright as in the visible spectrum, it is still highly reflective. A pixel with low NIR reflectance may correspond to liquid water (which absorbs NIR almost completely) or dense vegetation.

Additionally, the minimum variant of the Normalized Difference Snow Index (NDSI) is used. It is calculated as $(\text{Green} - \text{SWIR1}) / (\text{Green} + \text{SWIR1})$. An NDSI value greater than or equal to 100 (likely scaled) indicates a classical snow spectral signature: high reflectance in the green band and strong absorption in SWIR.

Regarding the snow fraction (SMA snow fraction): it must be greater than or equal to 20%. This means that at least one-fifth of the pixel must be composed of the snow/ice endmember.

For the slope variable, it must be greater than or equal to 0. This is a logical condition to discard NoData pixels in the DEM, ensuring that the evaluation only occurs in areas with defined terrain.

In summary, a pixel is classified as “Snow/Ice” only if it strictly meets the traditional spectral index criteria and exhibits a minimum proportion of snow in the physical mixture model. This dual validation (index + fraction) minimizes confusion with clouds, which are also bright in the visible spectrum but typically do not exhibit the same SWIR absorption signature nor the same spectral fraction.

In other words, if the cloud fraction computed by the SMA model exceeds the defined threshold, the pixel is reclassified as “Cloud” and masked out. This is essential in tropical regions, where clouds are the main contaminant in snow-cover time series.

4.3. Classification Variables

The classification algorithm uses several variables, organized into three main categories that capture different aspects of the spectral signature of snow-covered surfaces.

The first category comprises six spectral bands in their median wet-characteristic values (75th percentile of the NDSI).

The second category includes the use of spectral indices, highlighting the temporal variations of the Normalized Difference Snow Index (NDSI), calculated at three distinct moments:

- **Ndsi_median_dry:** Median NDSI for the dry-season or dry-characteristic mosaic (25th percentile), capturing the spectral condition of the permanent glacier without seasonal snow, or with minimal snow presence.
- **Ndsi_median_wet:** Median NDSI for the wet-season or wet-characteristic mosaic (75th percentile), reflecting the maximum snow cover during the wet period.
- **Ndsi_min:** Minimum NDSI value recorded throughout the year. A restrictive and relevant threshold for classification.

In snow-cover classification, topographic slope is not used as a direct discriminant variable, since the spatial distribution of seasonal snow may differ from the morphological pattern of the permanent glacier. Snow can accumulate on low-slope surfaces, meaning that a slope threshold could exclude valid snow-covered areas.

Finally, variables derived from Spectral Mixture Analysis (SMA) are included. SMA is a processing technique that decomposes the reflectance of each pixel into fractions of different pure materials, providing quantitative information about the sub-pixel composition of each area. The snow fraction is particularly relevant for distinguishing snow-covered surfaces from other land-cover types in high-mountain environments, while the cloud and shadow fractions help identify and filter pixels with compromised observations.

Table 1: Spectral bands used for classification.

Type	Name	Formula	Description	Reductor ¹				
				Median	Median dry	Median wet	Mín	Máx
Band	blue	B1 (L5 y L7); B2 (L8)	Visible spectrum Blue			X		
	green	B2 (L5 y L7); B3 (L8)	Visible spectrum Green			X		
	red	B3 (L5 y L7); B4 (L8)	Visible spectrum Red			X		
	nir	B4 (L5 y L7); B5 (L8)	Near Infrared			X		
	swir1	B5 (L5 y L7); B6 (L8)	Shortwave Infrared 1			X		
	swir2	B7 (L5); B8 (L7); B7(L8)	Shortwave Infrared 2			X		

Table 2: Spectral indices used for classification.

Type	Name	Formula	Description	Reductor ¹				
				Median	Median dry	Median Wet	Mín	Máx
Indices	NDSI	$\frac{(green - SWIR)}{(green + SWIR)}$	Normalized Difference Snow Index			X		

Table 3: SMA used for classification


Type	Name	Formula	Description	Reductor ¹				
				Median	Median dry	Median Wet	Mín	Má x
Fraction	Cloud fraction	SMA	Cloud fraction			X		
	Snow fraction	SMA	Snow fraction			X		

5. Post-Classification

Given the nature of the classification model and the processing of an extensive time series, a chain of post-classification filters was implemented.

5.1. Gap Filling

The sequence of post-processing filters begins with the Gap Fill, whose objective is to correct information gaps present in the snow-cover classification time series. This process is essential in the analysis of remote-sensing time series, as it addresses one of the most recurrent challenges in the discipline: the presence of no-data pixels or compromised classifications resulting from cloud cover, topographic shadows, or inherent limitations of satellite acquisition.

**Figure 5:** Gap Fill filter

The algorithm implements a bidirectional approach that operates in two complementary phases to ensure completeness of the time series. During the first phase, the algorithm traverses the series from the initial year to the final year. Whenever it encounters a no-data pixel each year, it assigns the classified value from the immediately preceding year, propagating information forward in time. The second phase performs the reverse traversal, from the final year back to the initial year, and fills pixels that remained without data after the first phase by assigning them the value from the nearest subsequent year.

with valid information. This bidirectional strategy ensures that gaps persist only in pixels lacking valid observations throughout the entire time series, minimizing information gaps through temporal imputation of the nearest classified value, without generating new spectral observations.

Unlike unidirectional filling, the bidirectional approach better captures state changes in the snow-cover signal. If a pixel transitions from “non-snow” to “snow” in a specific year, the backward and forward filling from that transition point preserves the change, preventing erroneous propagation of a state beyond its actual temporal extent.

It is important to remember that snow detection using passive optical sensors (Landsat) is inherently limited by cloud cover, topographic shadows, and acquisition failures. In the tropical Andes, where cloudiness is persistent for much of the year, there may be years without valid data for specific pixels or even entire regions.

The method implemented in this algorithm assumes that snow/ice cover in glacial zones exhibits high temporal inertia at the annual scale, especially in the Andes, where glaciers respond slowly to interannual climatic variability.

The code or algorithm reorganizes individual annual classifications into a multiband image where each band represents a specific year. This structure allows temporal operations to be applied efficiently.

Regarding methodological limitations, gap filling introduces some temporal uncertainty in the filled years. Snow-cover changes detected in years with real data are reliable, whereas filled years represent a conservative estimate based on temporal persistence. This limitation should be considered when interpreting interannual change rates.

Products with gap fill are ideal for long-term trend analysis (>5 years) and for continuous visualization of time series. For precise interannual change analysis, it is recommended to consult the products without gap fill, where gaps are explicitly maintained as “no data.”

5.2. Spatial filter

The spatial filter operates exclusively on the spatial dimension of the images, analyzing the coherence and connectivity of the pixels classified in each individual study year.

This filter addresses a common issue in satellite-image classification: the presence of isolated pixels or small groups of pixels classified differently from their immediate surroundings. These isolated pixels often represent classification errors rather than actual terrain features.

The spatial-filtering process is implemented in two main stages, applied independently for each year in the time series.

During the first stage, the filter uses the `connectedPixelCount` function to calculate how many glacier pixels are spatially connected. The second stage groups pixels and flags

all clusters with fewer than nine connected pixels. Additionally, the complementary focal-mode filter replaces isolated pixels with the predominant class value in their local neighborhood. Small clusters of isolated pixels are candidates for removal as potential errors.

For snow mapping in the Andes—where snow surfaces can have complex shapes due to mountainous topography—the spatial filter is valuable because it helps eliminate misclassifications caused by topographic shadows and small undetected clouds.

This entire strategy is grounded in the principle of spatial autocorrelation in geography: real geographic objects (such as glaciers and snowfields) tend to be spatially coherent and compact, whereas classification noise is spatially random and incoherent.

Operationally, an 8-neighbor connectivity (3×3 pixel window) is used to determine whether pixels are spatially connected.

On glacial surfaces, snow/ice pixels form continuous patches. Eight-neighbor connectivity (including diagonals) is more appropriate than four-neighbor connectivity because ice surfaces do not necessarily follow cardinal orientations and may extend diagonally across the pixel grid. This is particularly relevant for mountain glaciers with irregular geometries.

When discussing spatial connectivity, it is necessary to define the minimum threshold for connected-group size. Pixel groups with nine or fewer connected pixels are considered potential noise and are eligible for removal.

At 30-meter resolution, nine connected pixels represent an area of approximately 8,100 m² (0.81 hectares). Permanent glaciers and snowfields in the equatorial Andes rarely exist as patches smaller than 1 hectare, especially in accumulation zones.

The algorithm implements a temporal loop that applies spatial filtering independently for each year in the time series (1985–2025). Illumination conditions, cloud cover, and glacier state vary between years. Applying the spatial filter independently allows:

- Adaptation to the specific conditions of each year (e.g., years with more cloudiness may require more aggressive filtering),
- Preservation of real temporal changes in glacier and snow cover. A combined spatiotemporal filter could artificially smooth transitions between years where real changes occur.

In this and other filters, the code includes visualization layers to compare the original classification with the spatially filtered version. Visual validation is a critical tool for calibrating the `min_connect_pixel` parameter. A value that is too low leaves residual noise; a value that is too high removes legitimate snow/ice patches.




Figure 6: Spatial filter effect

5.3. Persistence Filter

The annual classifications generated in the previous sections identify the presence of snow/ice in each individual year. However, in the tropical Andes, a critical distinction exists between:

- Seasonal or ephemeral snow: Temporary snowfall that appears during the wet season and disappears during the dry season. These do not constitute permanent glaciers.
- Permanent glacier cover: Ice masses that persist across multiple years, regardless of seasonal variability.

The persistence filter evaluates the temporal frequency with which each pixel is classified as snow/ice throughout the entire time series. This metric allows discrimination between:

Persistent pixels: Those consistently classified as snow in a significant proportion of years, indicating stable glacier cover.

Sporadic pixels: Those classified as snow only occasionally, suggesting seasonal snow, classification errors, or transient processes.

Glaciers in the tropical Andes respond to climatic conditions with characteristic temporal inertia. Ideally, a healthy glacier maintains its ice cover year after year, even during adverse climatic events (e.g., El Niño years). In contrast, seasonal snowfall is highly sensitive to interannual climatic variability and disappears during dry years.

The use of a 35% threshold reflects the fact that Ecuador experiences high cloudiness throughout the year. Even with the previously applied gap-fill and spatial filters, some years may contain misclassifications due to residual atmospheric contamination.

Ecuadorian glaciers are undergoing accelerated retreat. Many small glaciers have disappeared during the study period (1985–2025). A 35% threshold captures glaciers that existed only during the early years of the series, while excluding seasonal snow that appears sporadically.

The algorithm iterates over all annual bands and accumulates a counter (snow_count) that records how many years each pixel was classified as snow.

Not all pixels have data for all years. After gap filling and spatial filtering, some pixels may remain masked (no data) in specific years. The denominator must reflect the years with real data for each pixel, not the theoretical total number of years. This avoids bias in pixels with incomplete time series.

Regarding limitations, the 35% threshold is an empirical assessment. Higher thresholds (>50%) exclude retreating glaciers that existed only partially during the period; lower thresholds (<20%) include persistent seasonal snow.

The filter assumes that the study period (1985–2025) is sufficiently long to capture climatic variability. For highly dynamic glaciers (rapid retreat or collapse), the filter may not adequately reflect recent changes.




Figure 7: Persistence Filter

5.4. Consistency filter

After applying the persistence filter, some pixels representing stable and healthy glaciers may still exhibit temporary gaps in their classification time series. These gaps occur when:

Spectral-classification failures: Years with particularly adverse atmospheric conditions (thin cirrus clouds, volcanic aerosols, extreme topographic shadows) that preprocessing could not fully remove.

- Limitations of temporal gap filling: The bidirectional gap-fill method implemented in Section 5.1. may not adequately capture glaciers with complex geometries or located in transition zones.
- Variability in reflectance: Changes in snow reflectance due to intense melt events, volcanic-dust deposition, or the presence of cryoconite, which temporarily alter the spectral signature.

As a solution to these conditions, the consistency filter identifies pixels classified as snow in a very high proportion of years ($\geq 96\%$) and assumes that these pixels represent unequivocally stable glaciers. For these pixels, the algorithm forces the snow classification in all years of the time series, filling occasional gaps that may exist.

The fundamental difference from the persistence filter is that persistence removes sporadic snow in non-persistent pixels—its objective is to distinguish glaciers from seasonal snow (persistence $\geq 35\%$)—whereas the consistency filter fills gaps in highly consistent pixels to ensure continuity in stable glaciers (consistency $\geq 96\%$).

Thus, the persistence filter acts as a low-pass filter (removing noise), while the consistency filter acts as a high-pass filter (reinforcing signal). Together, they produce a more robust and continuous time series.

5.5. Corrective filter

Despite the previous filters (gap fill, spatial, persistence, and consistency), anomalous years may still occur in which the snow/ice classification systematically fails across large areas that should be classified as glacier. These anomalous years may be caused by:

- Extreme atmospheric events: Volcanic eruptions depositing ash on glaciers, drastically altering their reflectance and temporarily darkening them.
- Exceptional climatic conditions: Years with persistent cloud cover that prevent the acquisition of cloud-free Landsat imagery over glacial zones.
- Preprocessing failures: Errors in the generation of annual mosaics affecting entire regions.

The filter uses an auxiliary layer of pixels persistent in 90% of the time series, representing the core of glacier/snow cover present in 37 of the 41 analyzed years (90% of $41 \approx 37$). For each year and region, the algorithm evaluates the relationship between the pixels in the 90% persistence layer and the pixels classified as snow cover in that year within the persistence zone. If the pixels in the 90% reference layer exceed the classified pixels in the evaluated year by more than 30%, that year is declared a gap year.

This condition is valid because pixels in the 90% layer should be contained within the annual 35% classification. When this containment relationship is reversed—i.e., when the permanent core exceeds the annual classification—it indicates the absence of a valid mosaic or deficient classification for that year.

For each year declared as a gap, the algorithm searches for the nearest valid years before and after it in the time series. The correction is applied using the intersection (logical AND operation) of the classifications from both neighboring valid years: only pixels identified as snow cover in both years are used to fill the gap year. This conservative criterion ensures that only the minimum glacier/snow cover verified through temporal continuity is used, avoiding overestimation.

6. Discussion and Data Limitations

Generating a snow/ice product for Ecuador requires a deep reflection on the very nature of the phenomenon being mapped. Unlike polar regions or mid-latitude mountains, where seasonal snow cover can persist for months, snow in the equatorial Andes is a meteorologically sporadic and thermodynamically unstable phenomenon.

Snowfall on Ecuadorian summits occurs under very specific atmospheric conditions, typically associated with the intrusion of cold air masses from the southern continent or extreme convective events over the Amazon. However, once deposited, snow at these latitudes is subjected to an unforgiving thermal regime, especially during daytime. Due to the equatorial position, incoming solar radiation is high year-round, and the 0 °C isotherm generally lies above 4,800–5,000 m a.s.l., meaning that fresh snow rarely persists for more than a few hours or days. This rapid melting is not an anomaly but the norm of the equatorial climate system.

The product described in this ATBD, based on annual median reflectance mosaics, is not designed to capture these ephemeral events. In fact, its methodological construction is deliberately oriented toward excluding them. The 35% persistence filter and the 96% consistency filter act precisely to remove the signal of very weak, sporadic snowfall and retain only snow cover that persists longer in time. Therefore, the resulting maps should be interpreted as a representation of snow cover undergoing transformation, not as an inventory of snowfall events.

This distinction is crucial for end users (hydrologists, glaciologists, and decision-makers). If the goal is to study snowfall frequency or the immediate hydrological contribution of fresh snow to daily melt, this product is a useful starting point, but other satellite inputs with higher spatial and temporal resolution should be explored. However, if the objective is to quantify long-term glacier retreat, the detection of seasonal snow is a methodological advantage, as it prevents temporary false positives from artificially inflating glacier-area estimates.

It is important to recall a key aspect that conditions interpretation of this product: the revisit frequency and spatial resolution of the Landsat constellation. Although the

30-meter resolution is excellent for mapping individual glaciers, Landsat's ability to observe the Earth's surface is limited by two factors that converge in the tropical Andes:

- Paramo ecosystems and Andean summits are subject to near-permanent convective cloud cover, especially in the afternoon. The methodology mitigates this through median mosaics and an internal cloud filter but cannot fully eliminate contamination from thin cirrus clouds or high-mountain haze.
- Landsat 8/9 has a 16-day revisit cycle (8 days combined). A snowfall that falls on a glacier and melts within 3 days will only be captured if the satellite happens to pass during that brief window. If the snowfall occurs on a cloudy day or between satellite passes, it will not be recorded.

This limitation reinforces the conclusion above: the product should not be used for analyzing seasonal dynamics of fresh snow or solid-precipitation events. The probability that Landsat captures ephemeral snow from a specific event is low, and the subsequent processing (medians + persistence filters) is explicitly designed to suppress any signal that does appear.

Additionally, this limitation introduces residual uncertainty in transition pixels. In years when a glacier is covered by fresh snow at the moment of satellite overpass, reflectance may temporarily increase, causing a retreating ice pixel (classified as non-snow in other years) to appear as snow that year. The consistency and corrective filters reduce this effect but do not eliminate it entirely.

Despite these limitations (and largely thanks to how the technical team explicitly addressed them in the processing chain) the snow/ice product generated by MapBiomás Agua constitutes a highly valuable scientific tool for monitoring climate change in the Andes.

Data availability is important because glacier retreat is one of the most sensitive and unequivocal indicators of global warming. Ecuadorian glaciers, in particular, are climate sentinels: their mass loss is rapid, quantifiable, and has direct consequences for water availability in paramo ecosystems and cities. A consistent 41-year time series (1985–2025), generated with a replicable and transparent methodology, allows precise quantification of how much ice has been lost and at what rate, while first discriminating seasonal snow.

Additionally, the data are important because the spatial and temporal continuity achieved through the gap-fill, persistence, consistency, and corrective filters is difficult to obtain with other global products (such as MODIS or VIIRS), which sacrifice spatial resolution for temporal resolution. This product fills a fine-scale data gap for a region where glaciers are small and require Landsat's 30-meter resolution to be mapped accurately.

Finally, it is essential that users interpret this product as a mask of maximum annual snow and ice cover, not as a precipitation record. Trends in snow-cover area derived from these maps are robust and comparable across years; however, minor interannual fluctuations (e.g., an apparent increase in snow during a wet year) should be analyzed

with caution. As discussed, such variations may reflect the greater or lesser likelihood that the satellite captured fresh snow, rather than real changes in underlying ice mass.

The processing chain explicitly acknowledges the climatic and logistical particularities of snow monitoring in Ecuador. By designing filters to remove ephemeral signals and reinforce persistent ones, and by documenting the inherent limitations of Landsat, the final product offers a conservative yet reliable view of snow cover—appropriate for the long-term studies demanded by the current climate crisis.

7. Conclusions

When referring to the spectral-classification strategy and reflectance thresholds, the implementation of sensor-specific reflectance thresholds proved to be an effective mechanism for the preliminary identification of snow and ice in the context of the Ecuadorian Andes. The differentiated parameterization between Landsat 5/7 and Landsat 8/9 helped mitigate the radiometric discrepancies inherent to each platform, ensuring temporal coherence across the 1985–2025 series. The integration of Spectral Mixture Analysis (SMA) as a complement to spectral thresholds added a physical dimension to the classification, enabling discrimination of mixed surfaces that traditional indices tend to overlook. The results suggest that combining threshold logic with mixture fractions increases classifier sensitivity in glacier-rock transition zones and reduces the incidence of false positives associated with clouds or highly reflective surfaces.

Regarding the differentiation of cover types through persistence percentiles, applying temporal persistence (35%) and consistency (96%) thresholds to the multiyear series enabled a clear methodological distinction between permanent glacier cover and ephemeral seasonal snow. The 35% persistence threshold acted as a discriminating filter that preserved glaciers with moderate interannual variability, while the 96% threshold reinforced the continuity of the most stable glacier cores, correcting omissions in years with adverse atmospheric conditions. Additionally, generating a 90% reference layer (P90) provided a robust basis for detecting gap years, allowing systematic deviations in annual classification to be quantified. This tiered percentile approach proved particularly suitable for tropical glaciers, where high cloudiness and rapid climatic responses demand flexible yet glaciologically grounded criteria.

With respect to the effectiveness of temporal and spatial filters in the processing chain, the sequential integration of spatial filters (connectivity and focal mode), temporal filters (bidirectional gap fill), and corrective filters (interpolation from neighboring years) was essential for producing a scientifically consistent final product. Spatial filtering effectively removed the salt-and-pepper noise characteristic of threshold-based classifications without compromising the integrity of glacier patches larger than one hectare. Meanwhile, the corrective filter based on the P90 reference layer enabled the identification and repair of years with systematic classification deficits, preventing the propagation of artificial gaps in the time series. Exporting with pyramids based on statistical mode ensured preservation of the dominant class during visualization and spatial aggregation processes. Taken together, the processing line was designed to handle the inherent limitations of optical remote sensing in tropical high-mountain

regions, providing a replicable methodological foundation for multitemporal glacier monitoring in other Andean environments.

8. References

- Bomshoms Calvelo, M., Quispe Gutiérrez, N., & Quispe Vega, K. (2018). Estudio de la frecuencia de nevadas en el Perú. Servicio Nacional de Meteorología e Hidrología del Perú (SENAMHI).
- Bradley, R. S., Vuille, M., Diaz, H. F., & Vergara, W. (2006). Threats to water supplies in the tropical Andes. *Science*, 312(5781), 1755-1756.
- Brown, R. D., & Robinson, D. A. (2011). Northern Hemisphere Spring snow cover variability and change over 1922–2010 including an assessment of uncertainty. *The Cryosphere*, 5(1), 219–229.
- Dietz, A. J., Kuenzer, C., Gessner, U., & Dech, S. (2012). Remote sensing of snow – A review of available methods. *International Journal of Remote Sensing*, 33(13), 4094–4134.
- Gascoin, S., Grizonnet, M., Bouchet, M., Salgues, G., & Hagolle, O. (2019). Theia Snow collection: High-resolution operational snow cover maps from Sentinel-2 and Landsat-8 data. *Earth System Science Data*, 11(2), 493–514.
- Lamantia, K. A., Larocca, L. J., Thompson, L. G., & Mark, B. G. (2024). El Niño enhances snow-line rise and ice loss on the Quelccaya Ice Cap, Peru. *The Cryosphere*, 18(10), 4633-4644.
- Luo, Jianfeng, Dong, Chunyu, Lin, Kairong, Chen, Xiaohong, Zhao, Liqiang, & Menzel, Lucas. (2022). Mapping forest snow cover using optical remote sensing, machine learning, and time-lapse photography. *Remote Sensing of Environment*. <https://doi.org/10.1016/j.rse.2022.113017>
- Romero, C. E., Saavedra, F. A., & Celis-Plá, P. S. M. (2026). Remote sensing methods for snowline dynamics detection in the Andes: a critical review. *EGUsphere*, 2026, 1-43.
- SENAMHI. (2018). Estudio de la frecuencia de nevadas en el Perú. Servicio Nacional de Meteorología e Hidrología del Perú.
- Schaeffli, B. (2015). Projecting hydropower production under future climates: a guide for decision-makers and modelers based on a review of published studies. *WIREs Water*, 2(4), 311–340.
- Vuille, M., Carey, M., Huggel, C., Buytaert, W., Rabatel, A., Jacobsen, D., ... & Sicart, J. E. (2018). Rapid decline of snow and ice in the tropical Andes—Impacts, uncertainties and challenges ahead. *Earth-science reviews*, 176, 195-213.