snowman: an open-source R package for automated 30-m snow and ice cover mapping using the Landsat archive

Author:

Niittynen, Pekka[1,2]*

Affiliations:

[1]Water, Energy and Environmental Engineering research unit, University of Oulu, Oulu, Finland
[2]Department of Biological and Environmental Science, University of Jyväskylä, Jyväskylä, Finland

*Correspondence to Pekka Niittynen, pekka.niittynen@oulu.fi

This is a Preprint submitted to *EarthArXiv* and has not been peer reviewed! The manuscript has also been submitted to *The Cryosphere* for peer review.

Keywords: snow, ice, Arctic, tundra, alpine, cryosphere, Landsat

snowman: an open-source R package for automated 30-m snow and ice cover mapping using the Landsat archive

Pekka Niittynen[1,2]

[1] Water, Energy and Environmental Engineering research unit, University of Oulu, Oulu, Finland
[2] Department of Biological and Environmental Science, University of Jyväskylä, Jyväskylä, Finland

Correspondence to: Pekka Niittynen (pekka.niittynen@oulu.fi)

Abstract. Seasonal snow and ice cover are critical components of the cryosphere yet mapping their dynamics at ecologically relevant spatiotemporal scales remains challenging. Here I present *snowman*, an open-source R package and algorithm for automated mapping of snow and ice cover dynamics at 30-m resolution using Landsat satellite imagery (1982–present). The algorithm combines globally trained probabilistic Random Forest classifiers with pixel-wise generalised additive models to estimate snow phenology metrics—including snow cover duration, snowmelt timing, and new-snow onset—across any location on Earth, without requiring specialist expertise in remote sensing. Trained on 691,925 manually labelled points from 529 Landsat scenes across 49 globally distributed sites, the classifier achieved an overall accuracy of 96.3% on an independent 15,000-point test dataset, compared to 80.0% for traditional normalised difference snow index-based (NDSI) approaches. Critically, *snowman* retained up to 2.2 times more usable observations than NDSI methods across a cloud-prone mountain landscape, enabling more detailed estimation of the snow dynamics. At two Finnish weather stations, *snowman* estimated snow cover duration, snowmelt timing, and new-snow onset to within 3–11 days of multi-year station records. Snow phenology maps showed strong spatial correspondence with independent fine-scale satellite-borne snow classifications (Pearson r = 0.79–0.83) and a high-resolution microclimate dataset (r = 0.82). The *snowman* algorithm is fully automated and scalable from personal computers to high-performance computing environments and offers a reproducible tool for snow and ice monitoring in climate science, hydrology, and ecological research.

1. Introduction

Seasonal snow cover is the most extensive component of the global cryosphere and plays a critical role in cold-region processes and drives feedbacks in the global climate system (Bokhorst et al., 2016; Pulliainen et al., 2020; Rixen et al., 2022). In Arctic, Antarctic, and alpine regions, snow blankets the ground for most of the year, dominating tundra landscapes (Essery et al., 2020). Due to its high albedo and low thermal conductivity, snow strongly influences local microclimates and permafrost (Zhang, 2005). It also impacts hydrology, plant productivity, species distributions, and overall ecosystem function (Chang et al., 2014; Gamon et al., 2013; Li et al., 2016; Niittynen and Luoto, 2018). Seasonal ice cover is another crucial part of freshwater ecosystems in cold regions, limiting light penetration into water and gas exchange between water and the

atmosphere (Prowse et al., 2006). Snow and ice not only shape high-mountain and polar ecosystems but also affect societies and livelihoods in the tundra and beyond, serving as freshwater reservoirs, impacting logistics, and causing hazardous events (Eerkes-Medrano and Huntington, 2021; Ford et al., 2014; Shivanna, 2022; Viviroli et al., 2007).

One of the primary challenges in snow and ice monitoring and modelling is their high spatial and temporal heterogeneity across landscapes (Voglimacci-Stephanopoli et al., 2022; Rutter et al., 2014). This is especially true in tundra regions, where open terrain allows wind to redistribute snow freely. Combined with rough topography, this creates substantial local-scale variability in snow conditions (Dery et al., 2004; Musselman et al., 2015). This heterogeneity typically occurs at spatial scales of tens or hundreds of metres (Elder et al., 1998; Mendoza et al., 2020; Miller et al., 2022), meaning that coarse-scale imagery or sparsely distributed measurement points fail to capture the variability. Moreover, freshwater bodies are often small in polar and mountainous regions (Muster et al., 2013; Verpoorter et al., 2014) making spatially explicit cryosphere mapping at ecologically relevant scales a persistent challenge (Kankaanpää et al., 2018; Niittynen et al., 2018, 2020).

Traditional in-situ measurement approaches can provide detailed information about snow and ice properties but are limited in spatial coverage and can be dangerous and expensive, particularly in mountainous and remote regions (Derksen and Mudryk, 2023; Hofer and Horak, 2020; Kongoli et al., 2019). Remote sensing enables systematic, large-scale observation of snow and ice conditions, including global mapping of snow cover extent and snow water equivalent (Dietz et al., 2012; Gascoin et al., 2024). Many applications have been built upon dense time series of coarse-scale optical, microwave, or radar imagery, capturing broad temporal trends in snow conditions globally, but at spatial resolutions too coarse for most ecological applications (Bormann et al., 2018; Notarnicola, 2022; Pulliainen et al., 2020). Conversely, snow conditions have been mapped in recent years at very fine spatial resolutions using repeated imaging campaigns with sensors onboard unmanned aerial vehicles (UAVs) (Bhardwaj et al., 2016; Rauhala et al., 2023). However, UAV flights do not provide long and dense time series and are spatially restricted to single landscapes. Fine-scale satellite instruments such as Sentinel-1 and -2 or PlanetScope offer sufficient spatial resolution and relatively frequent revisiting times for snow monitoring (Belward and Skøien, 2015; Gascoin et al., 2024) but lack the temporal extent of several decades, hindering studies on, for example, snow cover responses to climate variability.

The Landsat programme is the sole source of continuous, medium-resolution (30 m) satellite imagery spanning from 1982 to the present (Cohen and Goward, 2004). Since then, five Landsat satellites have produced multispectral Earth observations with broadly compatible technical specifications. The most recent launch, Landsat 9, occurred in September 2021. Landsat 8 and 9 are currently operational with long-term continuation planned for the program (Masek et al., 2020). The full Landsat archive was made openly available to the public in 2008, significantly benefiting researchers, environmental managers, and policymakers worldwide (Wulder et al., 2012).

Each Landsat satellite has a revisiting time of 16 days at the equator, but the overlap of flight paths increases towards the poles, providing more frequent images in Arctic and Antarctic regions (Ju and Masek, 2016; Masek et al., 2020). The Landsat mission primarily provides optical data, making the mapping of snow cover duration its main snow-related application (Macander et al., 2015). However, data quality is dependent on the absence of clouds, and separating clouds from snow can be challenging due to their sometimes-similar spectral signatures (Stillinger et al., 2019; Zhu et al., 2015). Additionally, data quality can be reduced, for instance, by topographic shadows common in mountainous regions (Panchard et al., 2023). These factors introduce large temporal gaps into Landsat time series, particularly in cloud-prone regions and during polar night. Processing such data further requires expertise in satellite image analysis and remote sensing programming (Gascoin et al., 2024; Niittynen and Luoto, 2018). Consequently, Landsat data remain underutilised in snow research, despite their demonstrated value across many scientific applications (Heegaard, 2002; Niittynen and Luoto, 2018; Panchard et al., 2023; Poussin et al., 2025).

To address these challenges, I developed an automated workflow wrapped into an R package called *snowman*. This open-source software library utilises decades of open data from the Landsat mission to map snow and ice cover dynamics at any location on Earth at a 30-m resolution without requiring specialist knowledge of satellite image processing. The methodology balances accuracy and computational requirements, running on any personal computer or high-performance computing system with internet access and R software installed. It does not require any sign-ins or third-party programs and is wrapped within five main user-friendly functions.

2. Methodology

2.1 General workflow

The *snowman* R package and algorithm map snow and ice cover dynamics across diverse landscapes using the full Landsat satellite imagery archive. This package offers a streamlined workflow for extracting and processing satellite data, calculating geospatial predictors, classifying imagery, and computing snow variables such as snow cover duration (Figure 1). Snow dynamics are best estimated across a multi-year period specified by the user, given the spatiotemporal sparsity of Landsat data. The following sections detail the methodology and steps for using the *snowman* package.

The *snowman* workflow supports internal parallelisation and has been tested on both Windows laptops and Linux-based supercomputers. It was developed using R versions 4.4.1 and 4.4.2 (R Core Team, 2025). All raster processing uses the *terra* R package (Hijmans, 2022), and spatial vectors are handled with the *sf* package (Pebesma, 2018; Pebesma and Bivand, 2023). The developmental version of the package is maintained on GitHub: https://github.com/poniitty/*snowman*

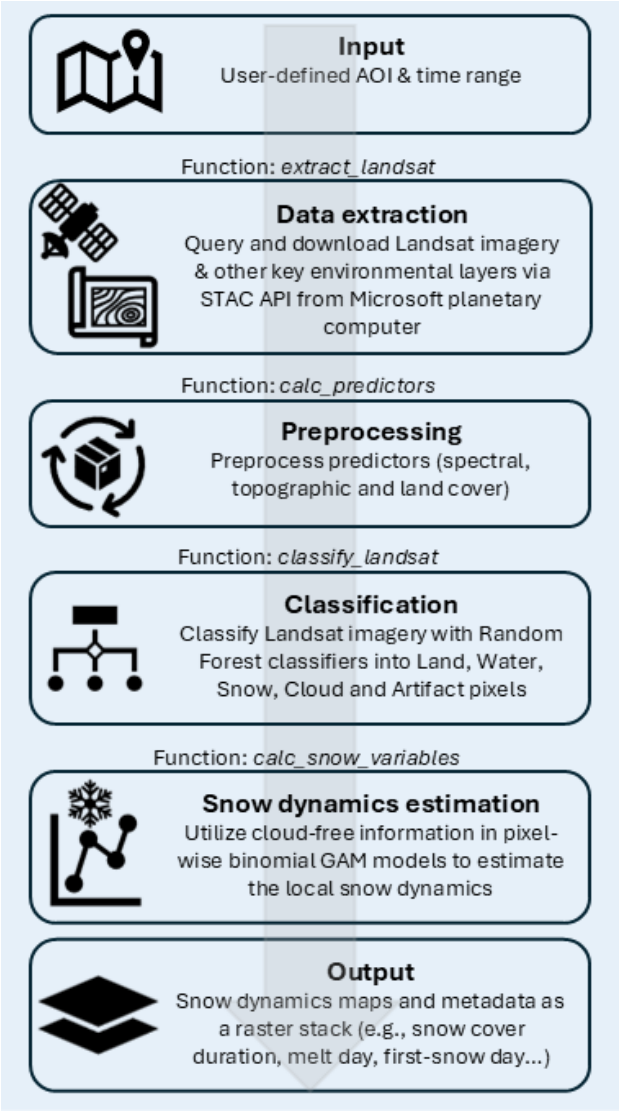

Figure 1. Workflow of the snowman package. AOI = area of interest; GAM = generalised additive model.

2.2 Raw data and data download

All data used to map the snow and ice dynamics are downloaded from Microsoft Planetary Computer (Microsoft Open Source et al., 2022) via the SpatioTemporal Asset Catalogues (STAC) API. STAC is a metadata standard that provides a unified structure for describing, cataloguing, and accessing spatiotemporal assets. STAC is utilised here with functions from the *rstac* package (Simoes et al., 2021). The Microsoft Planetary Computer is a cloud platform and geospatial data catalogue that uses open-source tools, supports open standards and includes all the source data products needed by the *snowman* algorithm.

The primary source of information for the snow and ice mapping is the multispectral Landsat imagery from Landsat TM4, Landsat TM5, Landsat ETM+7, Landsat OLI8 and Landsat OLI9 satellites (Earth Resources Observation and Science (EROS) Center, 2020a, b, c). The algorithm uses Landsat Collection-2 Level-2 Science Products, which include atmospherically corrected surface reflectance image data that have undergone initial quality filtering by the U.S. Geological Survey. Landsat data are available from 22 August 1982 to the present. All Landsat images are accompanied by a Quality Assessment (QA) band — a bit-packed layer encoding information on surface, atmosphere, and sensor conditions that affect pixel usability. This layer is generated from the CFMask algorithm (Foga et al., 2017; Zhu et al., 2015).

The *snowman* algorithm downloads all available Landsat images that overlap with the user's area of interest (AOI) and fall within the specified period. Users can select which Landsat satellites to include and can further filter the imagery based on the estimated scene cloud cover in the STAC metadata. Since the Microsoft Planetary Computer stores Landsat images in a cloud-optimised GeoTIFF format, the *terra* R package can access only the portion of the image overlapping the AOI, significantly reducing the data volume required for downloading. As part of the downloading process, the snowman algorithm ensures that all imagery uses the same grid system and is in the correct WGS84 UTM zone and projection.

After the images are downloaded, the algorithm reads all downloaded files, verifying data integrity. In case of problems, the algorithm tries to download the problematic images again. If the failure is repeated, the given image is removed from the rest of the analyses.

Next, the algorithm performs another layer of filtering by calculating the percentage of clear pixels (i.e., pixels that are neither cloud nor fill) in each Landsat image based on the QA band over the AOI. It removes all images from further analyses that have less than 10% of clear pixel coverage because very low clear-pixel fractions are often associated with reduced geolocation accuracy and a high proportion of difficult edge cases such as cloud boundaries and deep shadows.

Moreover, the 11-class WorldCover land cover classification (version 2021 v200) by the European Space Agency (ESA)(Zanaga et al., 2021) and ALOS World 3D-30m Global Digital Surface Model (DSM; version 3.2) by the Japan Aerospace Exploration Agency (Tadono et al., 2014) are downloaded from the Microsoft Planetary Computer and utilised in classifying the Landsat imagery.

2.3 Classifying Landsat imagery

The *snowman* algorithm uses pre-trained Random Forest classifier models to classify all non-missing pixels in each Landsat image into maps with five classes, each accompanied by probability maps. The following sections describe the training of these classifier models and their application to new imagery.

2.3.1 Training data

To develop a classifier model capable of distinguishing between snowy and icy surfaces, snow-free land (or ice-free water), and clouds, I manually collected training data. This involved visually interpreting five land cover classes from Landsat images and assigning labels to spatial points accordingly. I used QGIS geographic information system software (version 3.30.3)(QGIS Development Team, 2023) to visualise the Landsat images with multiple band combinations and subjectively selected training points across each target image. The five categories for the labels were:

1. Bare land, rock, or vegetated surface
2. Water
3. Snow or ice
4. Clouds
5. Artefacts

The artefact class represents clearly erroneous reflectance values, which sometimes occur in TM sensor imagery, particularly near the edges of the imaged area.

Training data were collected from 529 Landsat images across 49 areas (Supplementary Materials Figure S1). These training areas were selected from a set of 600 areas (each 20 km x 20 km), covering all major high-mountain and polar regions worldwide, which were gathered for a separate research project. The original 600 areas were partly selected subjectively (such as locations of active research stations or long-term measurement sites) and partly by random sampling within global tundra areas. Here, the sub-selection of training areas was iterative: Initially, I selected 20 areas representing the entire latitudinal gradient and various climatic and geological conditions. My prior field experience in many of these areas aided the visual interpretation of the imagery. For these areas, I randomly sampled 10 Landsat images and collected the first set of training data. Using this data, I fitted an initial classifier model and used it to classify a random set of Landsat images. I visually inspected the output and identified conditions where predictions were inaccurate, such as special types of clouds, deep shadows, glacial lakes, or very bright rock surfaces. Based on these evaluations, I selected additional training areas and Landsat imagery covering these challenging conditions, collected more training data from those, and repeated the modelling steps.

The training data intentionally emphasise land cover types and spectral conditions that are difficult to classify, rather than representing a random sample of all land covers. After several iterative rounds, no further systematic misclassification patterns were identified, indicating that the challenging conditions had been adequately represented in the training data. Thus, the final set consisted of 691,925 training points from 529 Landsat images across 49 areas.

2.3.2 Predictors

I gathered a set of predictors extracted for the training points. The predictors consist of the following categories:

a) **Raw surface reflectance values**: These include the raw surface reflectance values from the Landsat bands, excluding thermal bands, which provided no clear predictive benefit in initial tests and exhibited higher rates of artefacts and missing data.

b) **Reflectance differences to surroundings**: The difference between the reflectance of a focal cell and the mean reflectance of its neighbouring cells within a two-pixel buffer in the blue, NIR, and SWIR2 bands indicates whether the cell is exceptionally different from its surroundings. The three bands were selected to represent the main spectral characteristics of all bands.

c) **Local smoothed spectral statistics**: The minimum, maximum, and standard deviation in the blue, NIR, and SWIR2 bands, calculated with a two-pixel buffer around the focal cell. These statistics capture smoothed spectral signatures and help reduce local noise.

d) **Prior land cover class probabilities**: The probability of clouds, cirrus clouds, cloud shadow, snow, water, and clear conditions based on the QA band, calculated as the proportion of pixels classified into the class in question within a circular five-pixel buffer around the focal cell. In addition, the same calculation for clouds but with a larger 25-pixel buffer. These probabilities inform the model whether the surroundings of the focal pixel were predominantly assigned to a single class by the CFMask algorithm, potentially reducing the risk of misclassifying individual pixels within larger clouds.

e) **Reflectance differences to monthly median spectra**: This predictor captures how the reflectance of a pixel differs from the median (typical) reflectance for that month, using cloud-masked Landsat images from the same calendar period.

f) **Spatial heterogeneity differences**: The difference between the standard deviation calculated with a two-pixel buffer around the focal cell in the Landsat image and the standard deviation calculated with a two-pixel buffer around the focal cell in the monthly median raster. Only the NIR and SWIR2 bands were included, as they were considered sufficient to represent local spectral heterogeneity across all bands. This indicates whether local spectral heterogeneity differs substantially from typical conditions for that time of year.

g) **Normalised difference indices**: These indices are calculated over all possible pairs in the following list of spectral bands and predictors: blue, green, red, NIR, SWIR1, and SWIR2 from the Landsat image and from the spatial statistics described in (c) and (f) above.

h) **Mean spectral indices**: These are calculated over a set of cloud- and snow-masked Landsat images (including the three most cloud-free images from each month). Masking was based on the QA band. The spectral indices include (see the formulas of the indices in Supplementary Materials Text S1):

- Normalised Difference Vegetation Index (NDVI) (Rouse et al., 1974)
- Normalised Difference Snow Index (NDSI) (Riggs et al., 1994)
- Karst Bare-Rock Index (KBRI) (Pei et al., 2018)
- Sentinel Water Mask Index (SWM) (Milczarek et al., 2017)
- Landsat TM-based Brightness Index (BITM) (Mathieu et al., 1998)

These spectral indices were selected from the curated list by Montero et al. (2023), except for KBRI, which was created by Pei et al. (2018). These indices further inform the model about the expected land cover type of the pixel.

i) **Topographic predictors**: Based on the ALOS World 3D-30m DSM (Tadono et al., 2014), slope angle and hillshade with sun position matching the time of imaging the target Landsat image. These were calculated with the *terra* R package (Hijmans, 2022).

j) **Land cover classification**: The WorldCover 11-class land cover classification by ESA (Zanaga et al., 2021), mode-aggregated to the 30-m resolution from the original 10-m resolution. Moreover, the water proportion within the Landsat pixels is calculated as the proportion of water pixels in the 10-m resolution WorldCover data.

230

2.3.3 Fitting the classifier model

As the classifier models, I fitted fast, parallelised versions of Breiman's Random Forest models (Breiman, 2001), implemented in the *randomForestSRC* R package (Ishwaran et al., 2008; Ishwaran and Kogalur, 2025). Random Forest is a robust ensemble learning method widely used for land cover classification from satellite imagery (Kemppinen and Niittynen, 2022; Richiardi
235 et al., 2023). The method can handle high-dimensional data and model complex interactions across predictors (Breiman, 2001). I fitted the models separately for the TM, ETM+, and OLI sensors due to their slightly different spectral band specifications. Only the model for the TM sensor included class 5 (artefacts) in the response variable, as such clearly erroneous pixels did not occur in ETM+ or OLI imagery.

240 The number of candidate predictors for the classifier was high (134). Therefore, I conducted an initial round of modelling to determine the most influential predictors to be included in the final model. Reducing the number of predictors greatly decreases the computational time and memory requirements of the image classification. Here, I ran the initial models with all predictors included and identified the predictors most frequently used in tree-node splitting. I used the following model hyperparameters: splitrule = "auc", ntree = 100, nodedepth = 15, var.used = "all.trees", mtry = Inf, and nsplit = 100. After testing a range of
245 values in the number of included predictors (data not shown), I decided to select the 40 most frequently used predictors for fitting the final classifier. In fitting the final model, I experimented with a range of hyperparameter values for the number of trees and the minimum size of terminal nodes (data not shown) and ended up setting the number of trees to 30 for TM and OLI sensors and 40 for ETM+, with a node size of 10. Both parameters scale approximately linearly with computational time, and the final values represent a practical trade-off between processing speed and predictive accuracy.

250

The computation time for performing the classification increases with the amount of data. Thus, I ran initial tests to determine how much the data size affects the accuracy of the models (data not shown). In the end, I decided to use only a subset of 20% of the training points in the final models. This greatly shortened the time needed for predictions with only a minimal effect on model performance.

255

The final classifier models are then used to classify every pixel in all downloaded Landsat imagery and generate the probability values for each class.

2.4 Estimating snow dynamics

The final part of the process for calculating snow cover duration and other snow dynamic variables begins by calculating the proportion of different pixel classes (clouds, clear [land, water, and snow combined], land, and snow) covered in each Landsat image over the area of interest (AOI), based on the Random Forest classifications, as well as the proportion of missing data. Next, the *snowman* algorithm excludes all Landsat images with over 80% cloud cover because high scene cloud cover typically indicates lower geolocation accuracy and a high number of difficult cases for the classifiers, such as cloud edges and deep shadows. Processing continues only if at least 50 Landsat images remain, ensuring sufficient temporal coverage for reliable snow dynamics estimation.

Prior to the snow dynamics modelling, the classified imagery undergoes three preparatory steps. First, all pixels predicted to be clouds or artefacts, along with their immediate neighbouring pixels, are set to NoData. Second, the land and water classes are combined into a new "no-snow" class, and their probabilities are summed. Third, for the next steps, all snow and no-snow observations are included when the probability of the class is higher than 0.2. This allows both snow and no-snow observations from the same date to be retained, which is particularly useful during the melt season when pixels can contain a mixture of snow and bare ground. However, probability values below 0.2 were predominantly noise, and this threshold was selected to exclude the most uncertain classifier predictions. The class probability is utilised as the weight of the observations in the subsequent models, and therefore the less certain classifier predictions will have less influence on the output whatsoever.

In polar regions, the polar night prevents data collection by optical satellite sensors, and thus no Landsat data is available for this period. The *snowman* algorithm assumes that during the peak of winter, the polar regions are snow-covered by creating data for the missing winter period if needed. For each polar-night month lacking any Landsat imagery, synthetic snow observations are inserted on the 15th of that month, with the number of synthetic observations equal to the mean number of Landsat images available in the other months. This procedure is necessary to ensure that the pixel-wise GAM produces a realistic snow probability curve across the full annual cycle in polar regions. Users are encouraged to evaluate if this assumption holds in their polar study area and adjust the snow cover duration estimates accordingly if necessary.

The snowman algorithm estimates snow variables by iterating over all pixels one at a time. Snow dynamics are estimated with a generalised additive model (GAM) fitted to the classified pixel data using the *mgcv* package (Wood, 2017). The pixel-wise GAM results are provided only for pixels with at least 20 observations and where neither snow nor no-snow observations account for more than 99% of all classified observations. If these criteria are fulfilled, the binomial GAM is fitted with the following formula and syntax (Equation 1):

$$gam(snow \sim s(doy, bs = "cc", k = 5), knots = list(doy = c(1, 365)), family = "binomial", weights = prop, method = "REML")$$

Here, the binary snow/no-snow observations are treated as the response variable, related to the corresponding day of the year (DOY) of the observations, modelled with cyclic cubic regression spline smooths. To prevent overfitting and unrealistic snow phenology curves, the smooth term's basis dimension is limited to five. The knots argument informs the model of the start and end of the cycle (i.e., year). Furthermore, the model includes the probabilities that the Random Forest classifier predicted for the classifications as weights, giving more weight to more confident classification results. The model is fitted with the restricted maximum likelihood method.

The pixel-based GAM models are then used to predict snow probabilities for a sequence of days over a full year. The algorithm extracts the following 12 attributes from the pixel-wise data and GAM model predictions:

1. **Number of observations**: the number of unique dates with appropriate values in the Landsat data.
2. **Number of years**: the number of years with appropriate Landsat observations.
3. **Explained variance**: R^2 value of the GAM model, indicating how consistent the snow dynamics are.
4. **DOY with maximum snow probability**: the timing of peak winter.
5. **Maximum snow probability**: the highest snow probability.
6. **DOY with minimum snow probability**: i.e., when snow on the ground is least probable.
7. **Minimum snow probability**: the lowest snow probability.
8. **Raw snow cover duration**: the weighted mean of per-pixel snow/no-snow predictions (weighted by class probability) scaled to 0-365 days. This metric is derived directly from the Random Forest classifications, not from the GAM predictions.
9. **Snowy days**: the number of days in a year when GAM predicts snow probability > 0.5.
10. **Snow cover duration**: the sum of snow probabilities that GAM predicted for the sequence of days of a full year.
11. **Melting day of year**: the DOY when the probability of snow in GAM predictions first drops below 0.5 during the melting season. This will be NoData if the snow probability stays above or below 0.5 the entire year.
12. **New snow day of the year**: the DOY when the probability of snow first exceeds 0.5 during the autumn season. This will be NoData if the snow probability stays above or below 0.5 the entire year.

The output of the *snowman* algorithm is a raster stack that includes rasterised versions of all the 12 variables described above. It also includes three additional raster layers that can be beneficial for further analyses: the ESA WorldCover 11-class land cover classification mode-aggregated to the Landsat resolution; elevation from the ALOS World 3D-30m digital surface model; and the mean NDVI composite calculated over snow- and cloud-masked images (described above in the list of predictors for the Random Forest classifiers).

3. Usage example

The method and R package introduced here are applicable to any location on Earth. The full workflow took 37 minutes to process a 100 km² area in Finnish Lapland with 8 computing cores and 64 GB RAM in a Linux supercomputer with Intel Xeon processors running at 2.1 GHz for 2014-2023 with Landsat-8 imagery considered (138 used images; about 300 MB of downloaded data). The same example took 60 minutes by using 4 computing cores with a Lenovo Latitude 7430 laptop with an Intel i5-1245 processor (1.60 GHz) and 16 GB of RAM and by using a relatively slow internet connection (max 10 Mb/s).

The *snowman* R package is built around five main functions. A detailed tutorial and a user case with accompanying R code can be found in the Supplementary Materials (**Text S2**).

1. *extract_landsat*: extracts Landsat imagery for the specified AOI and period from the Microsoft Planetary Computer using the *rstac* R package (Simoes et al., 2021).
2. *calc_predictors*: calculates the predictors to be utilised in classifying the Landsat imagery with the Random Forest classifier model.
3. *download_model*: downloads the classifier model objects from a GitHub repository (https://github.com/poniitty/*snowman*_models).
4. *classify_landsat:* classifies the Landsat imagery using the downloaded models into snowy, snow-free, and cloudy pixels.
5. *calc_snow_variables*: calculates the snow variables across the AOI, which are the final products of the *snowman* algorithm. The output is a SpatRaster object with 15 raster layers (described in Sect. 2.4).

If the user wants to map the snow over a larger area, my recommendation is to split the whole AOI into subpolygons (e.g., with the *st_make_grid* function from the *sf* R package) and run the *snowman* over the subpolygons one at a time and combine the output at the end of the full process chain. Here, I illustrate the output of the *snowman* algorithm across the whole of Iceland (**Figure 2**). I divided the area of Iceland into 20 km x 20 km subpolygons and run snowman iteratively over them in a high-performance computing system.

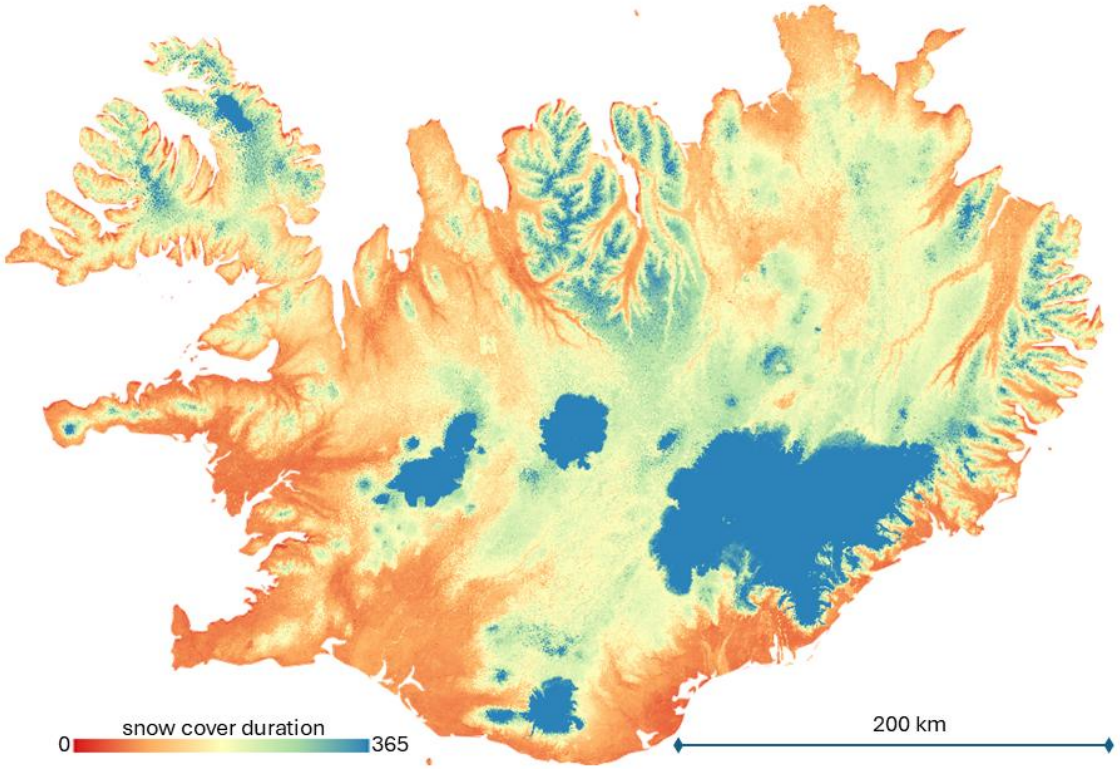

Figure 2. Snow and ice cover duration in Iceland in 2014-2023 mapped at 30-m resolution with the snowman algorithm.

4. Validation

To validate the output variables produced by the *snowman* algorithm, I compared the results to snow cover duration maps generated using two simpler Landsat-based methods, similar to those used by Rumpf et al. (2022) and Panchard et al. (2023). In these examples, I calculated snow cover duration and other related variables over the period 2014-2023 using Landsat OLI 8 imagery.

First, I used the QA band to mask all cloudy and cloud shadow pixels. For the masked images, I calculated the Normalised Difference Snow Index (NDSI, see Section 2.3.2) and binarised the resulting index values into snowy and no-snow pixels (NDSI >= 0.4 is considered snow if the reflectance of the red and NIR bands is > 0.1, which excludes water and deep shadows that can also have high NDSI values). Next, I used these binarised snow maps to calculate snow cover duration using two methods:

- **$NDSI_{raw}$ method**: In this first comparative method, I simply calculated an average over all snow/no-snow observations and multiplied the resulting proportion of snow observations by 365 to obtain a raw estimate of the number of snowy days in each pixel. For areas experiencing polar night, I added simulated observations for months without any data, similar to the approach used in the *snowman* algorithm described above.

- **$NDSI_{gam}$ method**: In this second comparative method, I used the same GAM-based method to estimate snow cover duration as in the *snowman* algorithm, except for not using weights in the model, as the NDSI-based binarisation does not produce probabilities for the snow or no-snow classes.

I selected three study areas to showcase and visualise the snow cover duration produced by either the *snowman*, $NDSI_{raw}$, or $NDSI_{gam}$ methods: one area in northern Finland (Kilpisjärvi), one in the Swiss Alps (Furka Pass), and one in Hokkaido, Japan (Mount Asahi).

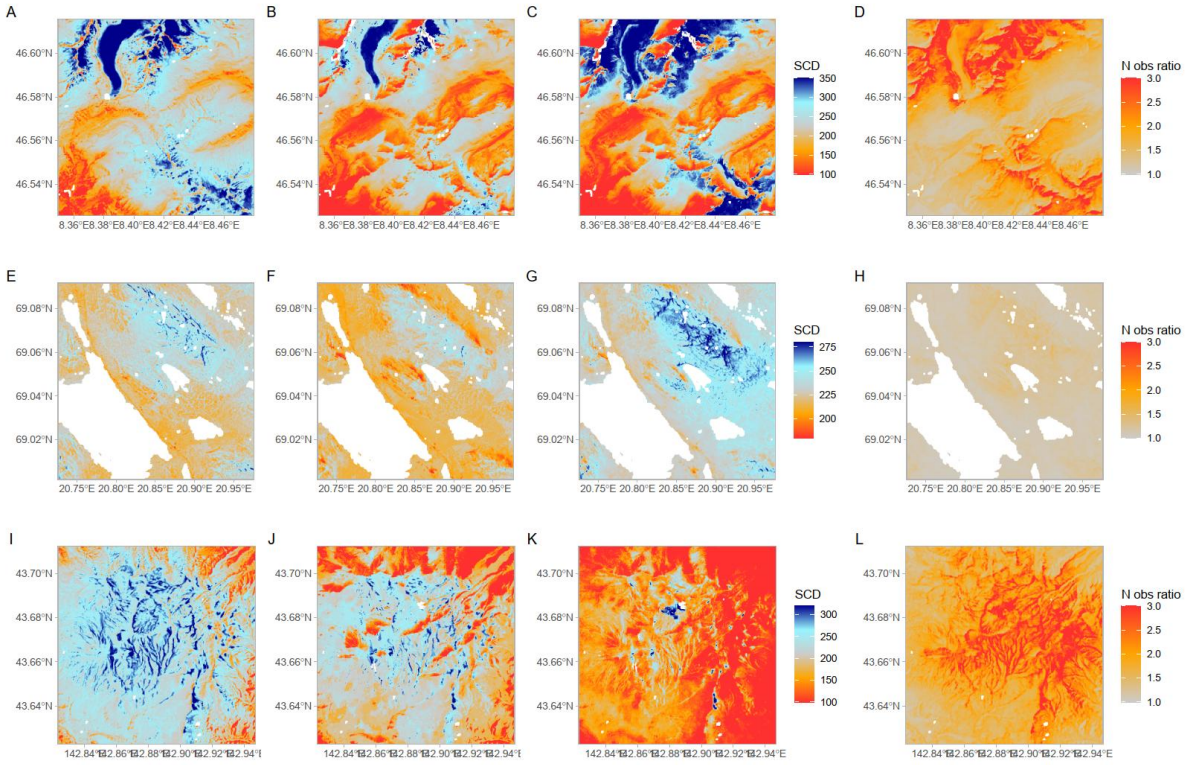

Figure 3. Comparison of three methods to calculate snow cover duration (SCD): snowman (A,E,I); $NDSI_{gam}$ (B,F,J); and $NDSI_{raw}$ (C,G,K) The methods were tested across three study areas: Furka Pass, Switzerland (A-D); Kilpisjärvi, Finland (E-H); Mount Asahi, Japan (I-L). Each area is 10 km x 10 km. The rightmost panels (D,H,L) show the ratio of the number of observations between snowman and QA band-based algorithms: N obs ratio = number of observations ratio. White areas represent water bodies.

In all three illustrated cases (**Figure 3**), the *snowman* algorithm clearly shows more fine-scale heterogeneity in snow cover duration than the NDSI-based methods. The largest differences occur relative to the simplest $NDSI_{raw}$ method, which loses most spatial detail and can substantially overestimate or underestimate snow cover duration. Notable differences also emerge at lower elevations in Japan, where *snowman* estimates longer snow cover duration beneath evergreen conifer canopies, conditions where NDSI characteristically underestimates snow. Furthermore, NDSI methods cannot detect snow under deep

topographic shadows, which should therefore be masked from the output, thus missing a lot of information that can be reliably used in the *snowman* algorithm.

At the Finnish study site, the *snowman* algorithm provides only slightly more usable observations per pixel (on average, 1.15 times more), but at the Swiss and Japanese sites, the difference is much higher in favour of the *snowman* algorithm (1.7 and 2.2 times more observations, respectively). Locally, the difference can be much higher; for example, at the Swiss site, the minimum number of observations in the NDSI methods is only 14, whereas it is 209 in the *snowman* algorithm. The greatest loss of potentially useful observations occurs near late-lying snow and ice, where the QA band often classifies clear observations erroneously as clouds.

At the Swiss site, the snow cover duration produced by the *snowman* algorithm has a Spearman correlation coefficient of 0.722 with the output of $NDSI_{gam}$ and 0.680 with $NDSI_{raw}$. The respective correlations are 0.659 and 0.686 in Finland and 0.803 and 0.806 in Japan. On average, the *snowman* algorithm estimated a 44-day longer snow season in Switzerland than $NDSI_{gam}$ (the mean absolute difference was 47 days), while in Finland, the average difference was 9 days (10 days), and in Japan, it was 39 days (40 days). These differences are largely attributable to *snowman*'s improved detection of snow in forested areas and around late-lying snow patches.

4.1 Validation of the Random Forest classifier

To assess the accuracy of the Random Forest classifier in classifying Landsat pixels, I randomly sampled a total of 150 individual Landsat images (50 images for each Landsat sensor type) across the 600 cold ecosystem study sites worldwide (see Section 2.3.1 for a description). From these 150 images, I took a random sample of 100 spatial points within each image frame, totalling 15,000 points. I visually inspected all these points and assigned the true class for each pixel in which the points fell. This dataset represents an independent test dataset that was not used in training the classifier and is largely from geographic areas outside the training data collection.

Following this step, I extracted the Random Forest classifier predictions for all the points. Additionally, I calculated the pixel classifications based on the QA band and NDSI method described above (i.e., clouds and water based on CFMASK classification and separating snow/ice from other clear pixels based on NDSI values). Here, I constructed NDSI-based classifications with two NDSI thresholds: 0.15 and 0.4, both used in previous literature (Macander et al., 2015; Panchard et al., 2023). Finally, I compared the *snowman* predictions and the two NDSI-based classifications to the manually labelled independent test dataset.

Overall, over 96% of the pixels in the test dataset were correctly classified with the *snowman* algorithm (**Table 1**), which was clearly higher than with the two NDSI-based methods (80%). The *snowman*'s Random Forest classifiers performed

consistently well across land cover types, under cloud shadows, and across all three Landsat sensor types. The NDSI-based methods achieved relatively good classification accuracy (~90%) under clear conditions (i.e., clouds and cloud shadows excluded first) but performed poorly over water bodies and under cloud shadows. Importantly, the *snowman* algorithm yields 42% more usable observations than NDSI-based methods when clouds and cloud shadows are excluded first. This difference would be even higher if topographic shadows were also excluded. The *snowman* algorithm and the NDSI methods are comparably good at detecting cloudy pixels, but the *snowman* algorithm performs clearly better in separating land, water, and snow in cold climate regions. The accuracy in detecting the clearly erroneous "artefact" pixels in Landsat data could not be evaluated, as only two artefact pixels ended up in the test data, indicating how rare these cases are.

Table 1. Evaluation of the pixel classification based on the independent test dataset of 15,000 observations. The proportion of correctly classified pixels with the snowman algorithm and two NDSI-based methods. Clear pixel includes all pixels that were not cloud or cloud shadow based on the CFMask classification in the QA bands of Landsat images.

	N. of points	*snowman*	$NDSI_{0.15}$	$NDSI_{0.40}$
All	15000	0.963	0.799	0.800
Only clear pixels (based on CFMask)	6913	0.985	0.896	0.899
Not cloud shadows (based on CFMask)	13403	0.963	0.827	0.828
Cloud shadows (based on CFMask)	1597	0.968	0.566	0.566
Over water (based on ESA WorldCover)	821	0.954	0.733	0.739
True class = "Land"	5851	0.986	0.742	0.750
True class = "Water"	356	0.938	0.694	0.694
True class = "Snow"	3433	0.961	0.709	0.700
True class = "Cloud"	5358	0.941	0.927	0.927
True class = "Artefact"	2	0.500	0.000	0.000
Landsat OLI (8 or 9)	5000	0.957	0.745	0.746
Landsat ETM+ (7)	5000	0.962	0.822	0.823
Landsat TM (4 or 5)	5000	0.971	0.830	0.833

4.2 Validation with weather station data

To assess how the snow cover duration produced by the *snowman* algorithm compares to observed snow cover duration at weather stations, I downloaded the full Global Historical Climatology Network - Daily (GHCN-Daily), Version 3 dataset.

Initially, 201 weather stations were found within the 600 study regions. However, a visual inspection revealed geolocation inaccuracies, for instance, with coordinates falling on lakes or ocean surfaces. Additionally, many weather stations were located close to buildings, roads, and other human infrastructure that could disturb the spectral signal and bias the snow estimation from the Landsat imagery. A comparison between *snowman* results and weather station-derived snow parameters requires accurate spatial overlay. Thus, I decided to rely only on two stations whose positions could be verified to within ±5 m using aerial imagery: Kilpisjärvi (Kyläkeskus) and Värriö (Värriötunturi) in northern Finland. I downloaded the snow depth data from the original source, the Finnish Meteorological Institute. I calculated the number of days with snow depth over zero per year and averaged it over the 2014-2023 period. I also used a similar binomial GAM model to estimate the average snowmelt and new-snow dates with binarised snow depth data, as used in the *snowman* algorithm with Landsat imagery.

At the Kilpisjärvi weather station, the weather data showed that the average count of snow days in 2014-2023 was 219 days. The average snowmelt day of the year (DOY) was 144, and the new-snow onset DOY was 287. Landsat imagery analysed with the *snowman* algorithm showed respective values of 218, 141, and 288, meaning that all estimates deviate by a maximum of 3 days from the weather station record. The weather station is in an open area but relatively close to settlements. However, snow conditions within the Landsat pixel footprint are unlikely to be substantially altered by nearby human activities.

At the Värriö weather station, the average count of snow days in 2014-2023 was 211 days. The average snowmelt DOY was estimated to be 139, and the snow onset DOY was 293. Landsat imagery analysed with the *snowman* algorithm showed respective values of 201, 128, and 293, thus underestimating the snow cover duration and snowmelt by approximately 10 days but matching the onset DOY. This springtime discrepancy may result from the local environment surrounding the weather station. The weather station is located within the courtyard of the Värriö research station, and the 30-m pixel includes portions of buildings and small roads where snow may be removed earlier. Additionally, the square where the weather station is located is bordered by tall pine trees, which have dark canopies in the spring when the snow on top of the canopy melts earlier than the snow on the ground. All these factors could disrupt the satellite snow signal, especially in the springtime, and lead to a slight underestimation of snow cover duration.

Furthermore, I used the daily snow depth data from the two weather stations to test how sensitive the snow cover duration estimation with the *snowman* algorithm is to the number of observations. I again binarised the weather station snow depth data into a snow/no-snow variable and repeatedly took random samples of the daily data with varying sample sizes. I then applied the GAM-based snow cover duration estimation to each subset. Random subsamples of 50, 100, 200, 300, 400, 500, 600, 700, and 800 observations were drawn, with 1,000 repetitions with each sample size, to quantify uncertainty in GAM-derived snow cover duration as a function of data density. The results of this simulation are shown in **Figure 4**.

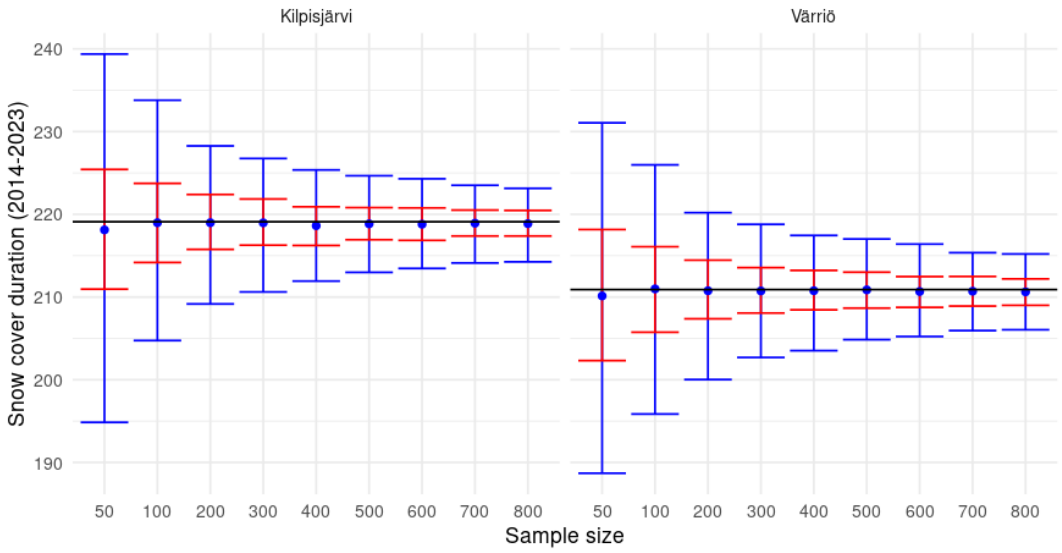

Figure 4. Sensitivity of the snow cover duration estimation in the snowman algorithm to sample size. Blue whiskers represent the 95% confidence intervals from 1000 repetitions, while red whiskers indicate the 50% confidence intervals. Blue dots denote the mean of the 1000 repetitions, and the black line represents the true snow cover duration calculated using the full weather station record.

4.3 Validation with a microclimate dataset

I compared the *snowman* algorithm's snowmelt estimates in Kilpisjärvi, Finland, to a local microclimate dataset described in Niittynen et al. (2024). Specifically, I used the start of season (SOS) map over the tundra landscape (excluding water bodies) produced at a 3-m spatial resolution with a statistical model based on hundreds of microclimate logger time series and fine-scale remote sensing. I mean-aggregated the SOS map to the same resolution as the snowmelt map and calculated the Pearson correlation coefficient (R) across the overlapping pixels. The correlation coefficient was 0.82. I did not calculate the difference between the variables due to differing time windows and because SOS does not directly equate to the snowmelt date. In regions where snow is blown away throughout the winter, the snow disappears much earlier than the thermal growing season starts. Furthermore, a known shortcoming of the microclimate dataset is that the original logger data did not cover the latest melting snowbed environments and thus consistently underestimates the local SOS in those habitats. Due to these discrepancies, I consider the shown correlation to be high.

4.4 Validation with fine-scale PlanetScope imagery

PlanetScope satellites (2014-present) provide multispectral Earth observations with nearly daily temporal coverage at a 3-m resolution. I downloaded clear-sky PlanetScope images for three dates for the Kilpisjärvi study site, with each image taken a maximum of two days away from a corresponding Landsat-8 image. I utilised an algorithm described and used earlier in Kemppinen and Niittynen (2022), Rissanen et al. (2023), and Heikkinen et al. (2024) to classify the PlanetScope images into

snow, water, or land classes. This Random Forest-based algorithm is conceptually analogous to the *snowman* algorithm described here. Based on the snow/snow-free classification, I calculated the snow-covered proportion of each Landsat pixel and correlated this with the Landsat-based snow probability maps produced by the Random Forest classifier from the *snowman* algorithm.

The general patterns of the Landsat-based snow probability maps and the PlanetScope-derived fractional snow cover are very similar (**Figure 5**), with strong correlations across the three image pairs: Pearson correlation coefficients of 0.79 on 9th May 2018, 0.82 on 12-14th June 2017, and 0.83 on 23-25th July 2017. Notably, in the June 2017 images, snowmelt had clearly advanced during the two days between the imagery, which is the main reason the PlanetScope snow map shows visibly less snow than the Landsat image taken two days earlier. The relationship between snow probability and fractional snow cover was approximately linear, but the spread around the trend lines was relatively high, indicating that several other factors affect the class probabilities besides just the snow-covered area within the Landsat pixels. The classifier was not specifically trained to represent fractional snow cover, so the predicted probability should not be interpreted as a direct equivalent.

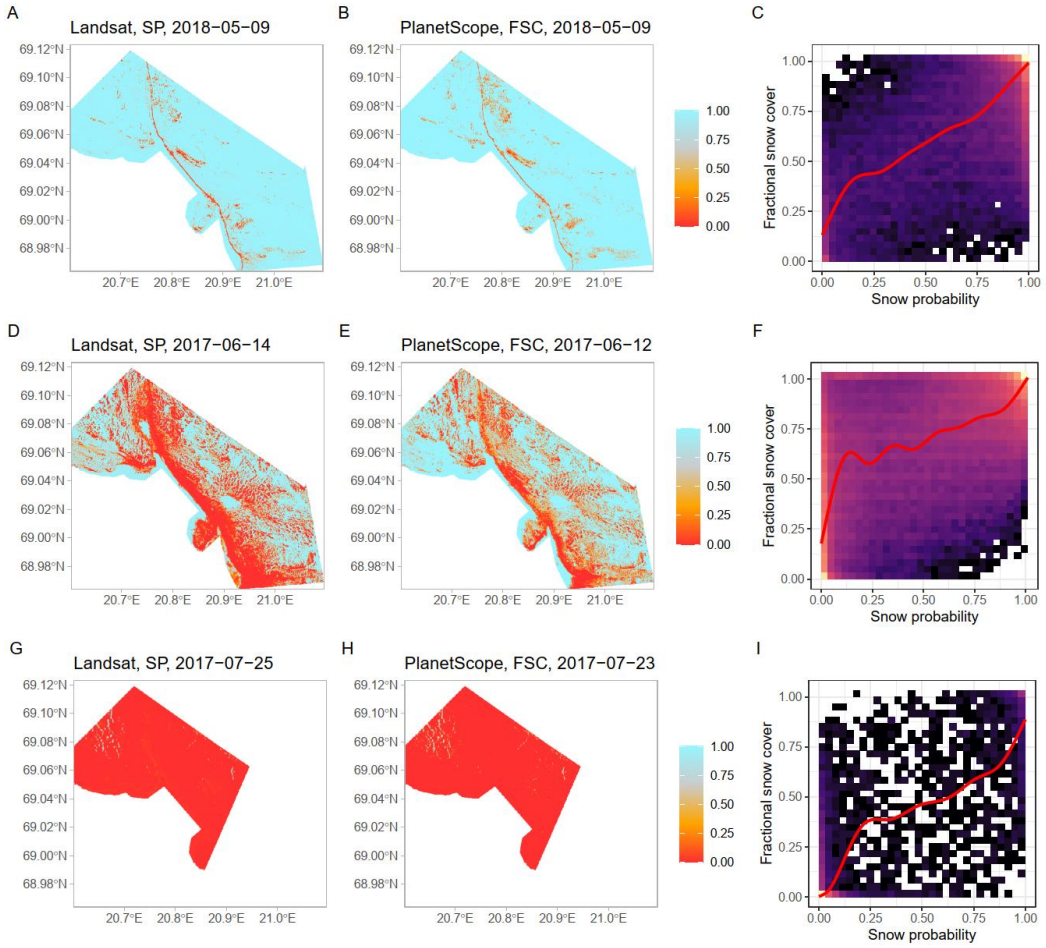

Figure 5. Comparison of snow cover estimates between three pairs of Landsat and PlanetScope images in Kilpisjärvi, northern Finland. The panels on the right-hand side (C, F, and I) show the heatmaps between the fractional snow cover from PlanetScope images and the snow probability from classifying the Landsat images with the snowman algorithm.

5. Discussion

5.1 Methodological advances

The *snowman* algorithm introduces key methodological advances in mapping snow and ice cover duration using Landsat imagery. Cloud-based platforms such as Google Earth Engine facilitate large-scale remote sensing analyses but impose constraints that limit algorithmic flexibility (Gorelick et al., 2017; Mateo-Garcia et al., 2018; Panchard et al., 2023; Rumpf et al., 2022). In contrast, *snowman* is fully open-source and designed for both local and high-performance computing systems, offering methodological transparency and extensibility often unachievable in cloud-based frameworks.

Historically, many Landsat-based snow cover products have relied on the CFMASK algorithm, designed for cloud and shadow masking rather than snow detection. Consequently, CFMASK often misclassifies persistent snow—especially small, late-lying patches or snow in shadows—as clouds. Similarly, methods based on the Normalised Difference Snow Index (NDSI) are widely used for their simplicity but are vulnerable to misclassification due to their reliance on fixed thresholds. For example, NDSI thresholds often fail in shadowed or forested areas where snow cover is obscured or spectral signatures are mixed.

The *snowman* algorithm addresses these limitations through several key innovations:

- **Probabilistic classification**: Instead of binary classification, the Random Forest model used in *snowman* outputs also class probabilities, providing nuanced information about uncertainty. This is especially valuable at snow patch edges and during transitional periods.
- **Improved vegetation handling**: The model better detects snow beneath canopies or among shrubs, where spectral signals are mixed. This increases the reliability of snow cover estimates in forested or shrub-dominated regions.
- **Shadow and thin cloud penetration**: The algorithm improves snow detection in deep topographic or cloud shadows and can distinguish snow through partially transparent thin clouds or mist and at cloud edges—conditions that commonly degrade other products.
- **Performance with lake ice**: unlike NDSI-based snow detection methods, the *snowman* algorithm is also trained to detect different ice conditions over water, thus enabling a variety of freshwater or coastal applications.
- **Enhanced observation density**: Validation shows that *snowman* can sometimes double the number of usable observations compared to QA-band and NDSI-based approaches. This feature is especially useful for studying snow cover dynamics over short periods or detecting long-term changes.
- **Automation**: The algorithm minimises user bias and enables reproducible, large-scale snow mapping without manual intervention. While more computationally intensive than threshold-based methods, the reproducibility and robustness of the workflow make it a valuable tool for researchers and practitioners working in cold-region environments.

5.2 Considerations for use and known limitations

Despite its strengths, several limitations and caveats must be considered when using the *snowman* package:

- **Temporal averaging and bias considerations**: Users should note that snowman outputs represent multi-year averages over the selected study period. In most areas, using several years of imagery is necessary to ensure sufficient data density and reliable estimates of snow dynamics. However, if the number of usable images differs greatly across years, the averages may be skewed toward periods with better coverage, particularly when combining imagery from multiple Landsat sensors with non-overlapping operational periods, which can unevenly weight certain years. To assess and mitigate such temporal bias, users are encouraged to examine the produced metadata to inspect the number of images per year. If significant disparities are present, a practical approach is to restrict analyses to a single Landsat sensor, ensuring a more balanced temporal sampling across the study period.
- **Limited coverage in early decades**: For studies focusing on snow dynamics during the 1980s and 1990s, users should consider the limited data availability in early Landsat records. In many regions, especially outside North America and Europe, there are notable gaps in image coverage during these decades, including periods of several consecutive years with no imagery. As with temporal bias, users should consult the provided metadata on image

counts to evaluate data completeness for the selected period. In data-sparse regions or periods, caution should be taken when interpreting results, and users may need to either extend the analysis period or focus on more data-rich decades.

- **Assumptions in polar night**: The algorithm assumes continuous snow cover during polar night months. While this is generally realistic over tundra and inland polar landscapes, it can lead to significant overestimation of ice cover duration over large lakes, fjords, or ocean areas that may remain ice-free long into the winter. Consequently, snow and ice cover estimates in coastal polar regions should be interpreted with caution or excluded.

- **Forest canopy effects**: Snow beneath dense evergreen conifers is difficult to detect, often leading to underestimation of snow cover duration. Users are advised to inspect the output and mask dense evergreen forest regions when the highest accuracy is required.

- **Spectral confounders**: Certain surfaces—such as bright calcareous rock, reflective rooftops, glacial lakes, and mixed water-land pixels—may be misclassified due to spectral similarities with snow or clouds. Users should visually inspect such areas and consider excluding or buffering problematic zones.

- **River and coastal pixels**: Rivers, especially large or rapidly flowing ones, and pixels along coastal boundaries can produce inconsistent classification results due to mixed land-water signals and wave-related spectral noise. These can sometimes be mixed with snow/ice and consequently lead to overestimated snow and ice cover.

- **Low sun angle and thin clouds**: In Arctic conditions during autumn, low solar angles and thin clouds over water can result in misclassification of water as snow or ice-free surfaces. Such cases should be interpreted with extra caution.

- **ETM+7 striping and path edge effects**: Landsat 7's scan line corrector failure introduces striping in observation count layers. Although *snowman* appears rather robust against this artefact—snow cover duration maps remain visually consistent—users should still be aware of potential observation biases when including Landsat ETM+7 imagery.

- **Bandwidth considerations**: The download size for Landsat imagery can be substantial, especially when long time series are requested over large areas. For example, a 20 × 20 km tile in Iceland may require 1.5 GB of imagery.

Despite these challenges, the validation analyses suggest that the *snowman* algorithm produces consistent and ecologically meaningful estimates of snow cover duration when averaged over multi-year periods. Importantly, the method's transparency and reproducibility support further customisation, allowing users to adapt the algorithm to their specific needs or data sources.

6. Conclusions

I present *snowman*, an open-source R package for reproducible, automated, fine-scale mapping of snow and ice cover duration and related cryospheric dynamics using the full Landsat surface reflectance archive. The package is applicable to any location on Earth using the full Landsat surface reflectance imagery archive. Unlike many conventional remote sensing approaches, *snowman* addresses known limitations in cloud masking, snow-cloud confusion, and threshold-based indices by employing a robust probabilistic classification algorithm and pixel-wise generalised additive models for snow phenology estimation. The method significantly improves snow cover monitoring in areas with complex terrain, frequent cloud cover, or topographic

shadows. Validation shows that *snowman* recovers up to twice the number of usable observations in cloud-prone environments compared to approaches relying on the Quality Assessment bands.

While computationally intensive, the algorithm is fully automated and scalable, making it suitable for both local desktop and high-performance computing environments. Looking ahead, *snowman* provides a foundation for future extensions, including the integration of Sentinel-2 data to improve spatiotemporal resolution (Gascoin et al., 2024) and data fusion approaches that combine optical remote sensing with climate and in-situ datasets for trend detection. Such developments would further support long-term analysis of snow and ice dynamics under climate change. By enabling fine-scale, reproducible snow monitoring across any landscape, *snowman* supports ecological and environmental research at the spatial scales most relevant to ecosystem processes and functioning.

Appendices

Appendix A: Supplementary Materials

- Supplementary Figure S1
- Supplementary Text S1
- Supplementary Text S2

Code, data, or code and data availability

The source code of the snowman R package is maintained and openly available on GitHub: https://github.com/poniitty/snowman

Author contributions

Pekka Niittynen is responsible for all method development, analyses and writing.

Competing interests

The author declares no competing interests.

Acknowledgements & Financial support

PN acknowledges funding from the Research Council of Finland (grant no. 378397; 347558; PROFI8: 365202), the Kone Foundation and the Nessling Foundation.

References

Belward, A. S. and Skøien, J. O.: Who launched what, when and why; trends in global land-cover observation capacity from civilian earth observation satellites, ISPRS J. Photogramm. Remote Sens., 103, 115–128, https://doi.org/10.1016/j.isprsjprs.2014.03.009, 2015.

Bhardwaj, A., Sam, L., Akanksha, Martin-Torres, F. J., and Kumar, R.: UAVs as remote sensing platform in glaciology: Present applications and future prospects, Remote Sens. Environ., 175, 196–204, https://doi.org/10.1016/j.rse.2015.12.029, 2016.

Bokhorst, S., Pedersen, S. H., Brucker, L., Anisimov, O., Bjerke, J. W., Brown, R. D., Ehrich, D., Essery, R. L. H., Heilig, A., Ingvander, S., Johansson, C., Johansson, M., Jonsdottir, I. S., Inga, N., Luojus, K., Macelloni, G., Mariash, H., McLennan, D., Rosqvist, G. N., Sato, A., Savela, H., Schneebeli, M., Sokolov, A., Sokratov, S. A., Terzago, S., Vikhamar-Schuler, D., Williamson, S., Qiu, Y. B., and Callaghan, T. V.: Changing Arctic snow cover: A review of recent developments and assessment of future needs for observations, modelling, and impacts, Ambio, 45, 516–537, https://doi.org/10.1007/s13280-016-0770-0, 2016.

Bormann, K. J., Brown, R. D., Derksen, C., and Painter, T. H.: Estimating snow-cover trends from space, Nat. Clim. Change, 8, 924–928, https://doi.org/10.1038/s41558-018-0318-3, 2018.

Breiman, L.: Random forests, Mach. Learn., 45, 5–32, https://doi.org/10.1023/a:1010933404324, 2001.

Chang, J., Wang, G., Gao, Y., and Wang, Y.: The influence of seasonal snow on soil thermal and water dynamics under different vegetation covers in a permafrost region, J. Mt. Sci., 11, 727–745, https://doi.org/10.1007/s11629-013-2893-0, 2014.

Cohen, W. B. and Goward, S. N.: Landsat's role in ecological applications of remote sensing, Bioscience, 54, 535–545, https://doi.org/10.1641/0006-3568(2004)054%5B0535:lrieao%5D2.0.co;2, 2004.

Derksen, C. and Mudryk, L.: Assessment of Arctic seasonal snow cover rates of change, The Cryosphere, 17, 1431–1443, https://doi.org/10.5194/tc-17-1431-2023, 2023.

Dery, S. J., Crow, W. T., Stieglitz, M., and Wood, E. F.: Modeling snow-cover heterogeneity over complex arctic terrain for regional and global climate models, J. Hydrometeorol., 5, 33–48, https://doi.org/10.1175/1525-7541(2004)005%3C0033:MSHOCA%3E2.0.CO;2, 2004.

Dietz, A. J., Kuenzer, C., Gessner, U., and Dech, S.: Remote sensing of snow – a review of available methods, Int. J. Remote Sens., 33, 4094–4134, https://doi.org/10.1080/01431161.2011.640964, 2012.

Earth Resources Observation and Science (EROS) Center: Landsat 4-5 Thematic Mapper Level-2, Collection 2 [dataset], , https://doi.org/10.5066/P9IAXOVV, 2020a.

Earth Resources Observation and Science (EROS) Center: Landsat 7 Enhanced Thematic Mapper Plus Level-2, Collection 2 [dataset], , https://doi.org/10.5066/P9C7I13B, 2020b.

Earth Resources Observation and Science (EROS) Center: Landsat 8-9 Operational Land Imager / Thermal Infrared Sensor Level-2, Collection 2 [dataset], , https://doi.org/10.5066/P9OGBGM6, 2020c.

Eerkes-Medrano, L. and Huntington, H. P.: Untold Stories: Indigenous Knowledge Beyond the Changing Arctic Cryosphere, Front. Clim., 3, https://doi.org/10.3389/fclim.2021.675805, 2021.

Elder, K., Rosenthal, W., and Davis, R. E.: Estimating the spatial distribution of snow water equivalence in a montane watershed, Hydrol. Process., 12, 1793–1808, https://doi.org/10.1002/(SICI)1099-1085(199808/09)12:10/11%3C1793::AID-HYP695%3E3.0.CO;2-K, 1998.

Essery, R., Kim, H., Wang, L., Bartlett, P., Boone, A., Brutel-Vuilmet, C., Burke, E., Cuntz, M., Decharme, B., Dutra, E., Fang, X., Gusev, Y., Hagemann, S., Haverd, V., Kontu, A., Krinner, G., Lafaysse, M., Lejeune, Y., Marke, T., Marks, D., Marty, C., Menard, C. B., Nasonova, O., Nitta, T., Pomeroy, J., Schädler, G., Semenov, V., Smirnova, T., Swenson, S., Turkov, D., Wever, N., and Yuan, H.: Snow cover duration trends observed at sites and predicted by multiple models, The Cryosphere, 14, 4687–4698, https://doi.org/10.5194/tc-14-4687-2020, 2020.

Foga, S., Scaramuzza, P. L., Guo, S., Zhu, Z., Dilley, R. D., Beckmann, T., Schmidt, G. L., Dwyer, J. L., Joseph Hughes, M., and Laue, B.: Cloud detection algorithm comparison and validation for operational Landsat data products, Remote Sens. Environ., 194, 379–390, https://doi.org/10.1016/j.rse.2017.03.026, 2017.

Ford, J. D., McDowell, G., and Jones, J.: The state of climate change adaptation in the Arctic, Environ. Res. Lett., 9, 104005, https://doi.org/10.1088/1748-9326/9/10/104005, 2014.

Gamon, J. A., Huemmrich, K. F., Stone, R. S., and Tweedie, C. E.: Spatial and temporal variation in primary productivity (NDVI) of coastal Alaskan tundra: Decreased vegetation growth following earlier snowmelt, Remote Sens. Environ., 129, 144–153, https://doi.org/10.1016/j.rse.2012.10.030, 2013.

Gascoin, S., Luojus, K., Nagler, T., Lievens, H., Masiokas, M., Jonas, T., Zheng, Z., and De Rosnay, P.: Remote sensing of mountain snow from space: status and recommendations, Front. Earth Sci., 12, https://doi.org/10.3389/feart.2024.1381323, 2024.

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., and Moore, R.: Google Earth Engine: Planetary-scale geospatial analysis for everyone, Remote Sens. Environ., 202, 18–27, https://doi.org/10.1016/j.rse.2017.06.031, 2017.

Heegaard, E.: A model of alpine species distribution in relation to snowmelt time and altitude, J. Veg. Sci., 13, 493–504, https://doi.org/10.1111/j.1654-1103.2002.tb02076.x, 2002.

Heikkinen, J. M., Niittynen, P., Soininen, J., and Pajunen, V.: Patterns and drivers for benthic algal biomass in sub-Arctic mountain ponds, Hydrobiologia, 851, 689–708, https://doi.org/10.1007/s10750-023-05368-3, 2024.

Hijmans, R. J.: terra: Spatial Data Analysis, 2022.

Hofer, M. and Horak, J.: Extending Limited In Situ Mountain Weather Observations to the Baseline Climate: A True Verification Case Study, Atmosphere, 11, 1256, https://doi.org/10.3390/atmos11111256, 2020.

Ishwaran, H. and Kogalur, U. B.: Fast Unified Random Forests for Survival, Regression, and Classification (RF-SRC), manual, 2025.

Ishwaran, H., Kogalur, U. B., Blackstone, E. H., and Lauer, M. S.: Random survival forests, Ann. Appl. Stat., 2, 841–860, https://doi.org/10.1214/08-AOAS169, 2008.

Ju, J. and Masek, J. G.: The vegetation greenness trend in Canada and US Alaska from 1984–2012 Landsat data, Remote Sens. Environ., 176, 1–16, https://doi.org/10.1016/j.rse.2016.01.001, 2016.

Kankaanpää, T., Skov, K., Abrego, N., Lund, M., Schmidt, N. M., and Roslin, T.: Spatiotemporal snowmelt patterns within a high Arctic landscape, with implications for flora and fauna, Arct. Antarct. Alp. Res., 50, e1415624, https://doi.org/10.1080/15230430.2017.1415624, 2018.

Kemppinen, J. and Niittynen, P.: Microclimate relationships of intraspecific trait variation in sub-Arctic plants, Oikos, 2022, e09507, https://doi.org/10.1111/oik.09507, 2022.

Kongoli, C., Key, J., and Smith, T. M.: Mapping of Snow Depth by Blending Satellite and In-Situ Data Using Two-Dimensional Optimal Interpolation—Application to AMSR2, Remote Sens., 11, 3049, https://doi.org/10.3390/rs11243049, 2019.

Li, W. B., Wu, J. B., Bai, E., Jin, C. J., Wang, A. Z., Yuan, F. H., and Guan, D. X.: Response of terrestrial carbon dynamics to snow cover change: A meta-analysis of experimental manipulation (II), Soil Biol. Biochem., 103, 388–393, https://doi.org/10.1016/j.soilbio.2016.09.017, 2016.

Macander, M. J., Swingley, C. S., Joly, K., and Raynolds, M. K.: Landsat-based snow persistence map for northwest Alaska, Remote Sens. Environ., 163, 23–31, https://doi.org/10.1016/j.rse.2015.02.028, 2015.

Masek, J. G., Wulder, M. A., Markham, B., McCorkel, J., Crawford, C. J., Storey, J., and Jenstrom, D. T.: Landsat 9: Empowering open science and applications through continuity, Remote Sens. Environ., 248, 111968, https://doi.org/10.1016/j.rse.2020.111968, 2020.

Mateo-Garcia, G., Gomez-Chova, L., Amoros-Lopez, J., Munoz-Mari, J., and Camps-Valls, G.: Multitemporal Cloud Masking in the Google Earth Engine, Remote Sens., 10, 18, https://doi.org/10.3390/rs10071079, 2018.

Mathieu, R., Pouget, M., Cervelle, B., and Escadafal, R.: Relationships between Satellite-Based Radiometric Indices Simulated Using Laboratory Reflectance Data and Typic Soil Color of an Arid Environment, Remote Sens. Environ., 66, 17–28, https://doi.org/10.1016/S0034-4257(98)00030-3, 1998.

Mendoza, P. A., Musselman, K. N., Revuelto, J., Deems, J. S., López-Moreno, J. I., and McPhee, J.: Interannual and Seasonal Variability of Snow Depth Scaling Behavior in a Subalpine Catchment, Water Resour. Res., 56, e2020WR027343, https://doi.org/10.1029/2020WR027343, 2020.

Microsoft Open Source, McFarland, M., Emanuele, R., Morris, D., and Augspurger, T.: microsoft/PlanetaryComputer: October 2022, , https://doi.org/10.5281/zenodo.7261897, 2022.

Milczarek, M., Robak, A., and Gadawska, A.: Sentinel Water Mask (SWM) - new index for water detection on Sentinel-2 images, Crisis Information Centre at Space Research Centre, Polish Academy of Sciences, 2017.

Miller, Z. S., Peitzsch, E. H., Sproles, E. A., Birkeland, K. W., and Palomaki, R. T.: Assessing the seasonal evolution of snow depth spatial variability and scaling in complex mountain terrain, The Cryosphere, 16, 4907–4930, https://doi.org/10.5194/tc-16-4907-2022, 2022.

Montero, D., Aybar, C., Mahecha, M. D., Martinuzzi, F., Söchting, M., and Wieneke, S.: A standardized catalogue of spectral indices to advance the use of remote sensing in Earth system research, Sci. Data, 10, 197, https://doi.org/10.1038/s41597-023-02096-0, 2023.

Musselman, K. N., Pomeroy, J. W., Essery, R. L. H., and Leroux, N.: Impact of windflow calculations on simulations of alpine snow accumulation, redistribution and ablation, Hydrol. Process., 29, 3983–3999, https://doi.org/10.1002/hyp.10595, 2015.

Muster, S., Heim, B., Abnizova, A., and Boike, J.: Water Body Distributions Across Scales: A Remote Sensing Based Comparison of Three Arctic Tundra Wetlands, Remote Sens., 5, 1498–1523, https://doi.org/10.3390/rs5041498, 2013.

Niittynen, P. and Luoto, M.: The importance of snow in species distribution models of arctic vegetation, Ecography, 41, 1024–1037, https://doi.org/10.1111/ecog.03348, 2018.

Niittynen, P., Heikkinen, R. K., and Luoto, M.: Snow cover is a neglected driver of Arctic biodiversity loss, Nat. Clim. Change, 8, 997–1001, https://doi.org/10.1038/s41558-018-0311-x, 2018.

Niittynen, P., Heikkinen, R. K., Aalto, J., Guisan, A., Kemppinen, J., and Luoto, M.: Fine-scale tundra vegetation patterns are strongly related to winter thermal conditions, Nat. Clim. Change, 10, 1143-U134, https://doi.org/10.1038/s41558-020-00916-4, 2020.

Niittynen, P., Salminen, H., Pena-Aguilera, P., Aalto, J., Alahuhta, J., Luoto, M., Maliniemi, T., Rissanen, T., Roslin, T., Tukiainen, H., and others: A Gridded Microclimate Dataset from a Sub-Arctic Biodiversity Hotspot in Finland, bioRxiv, 2024–03, 2024.

Notarnicola, C.: Overall negative trends for snow cover extent and duration in global mountain regions over 1982–2020, Sci. Rep., 12, 13731, https://doi.org/10.1038/s41598-022-16743-w, 2022.

Panchard, T., Broennimann, O., Gravey, M., Mariethoz, G., and Guisan, A.: Snow cover persistence as a useful predictor of alpine plant distributions, J. Biogeogr., 50, 1789–1802, https://doi.org/10.1111/jbi.14689, 2023.

Pebesma, E.: Simple Features for R: Standardized Support for Spatial Vector Data, R J., 10, 439–446, 2018.

Pebesma, E. and Bivand, R.: Spatial Data Science: With Applications in R, Chapman and Hall/CRC, New York, 314 pp., https://doi.org/10.1201/9780429459016, 2023.

Pei, J., Wang, L., Huang, N., Geng, J., Cao, J., and Niu, Z.: Analysis of Landsat-8 OLI Imagery for Estimating Exposed Bedrock Fractions in Typical Karst Regions of Southwest China Using a Karst Bare-Rock Index, Remote Sens., 10, 1321, https://doi.org/10.3390/rs10091321, 2018.

Poussin, C., Peduzzi, P., Chatenoux, B., and Giuliani, G.: A 37 years [1984–2021] Landsat/Sentinel-2 derived snow cover time-series for Switzerland, Sci. Data, 12, 632, https://doi.org/10.1038/s41597-025-04961-6, 2025.

Prowse, T. D., Wrona, F. J., Reist, J. D., Gibson, J. J., Hobbie, J. E., Lévesque, L. M. J., and Vincent, W. F.: Climate Change Effects on Hydroecology of Arctic Freshwater Ecosystems, AMBIO J. Hum. Environ., 35, 347–358, https://doi.org/10.1579/0044-7447(2006)35%5B347:CCEOHO%5D2.0.CO;2, 2006.

Pulliainen, J., Luojus, K., Derksen, C., Mudryk, L., Lemmetyinen, J., Salminen, M., Ikonen, J., Takala, M., Cohen, J., Smolander, T., and Norberg, J.: Patterns and trends of Northern Hemisphere snow mass from 1980 to 2018, Nature, 581, 294–298, https://doi.org/10.1038/s41586-020-2258-0, 2020.

QGIS Development Team: QGIS Geographic Information System, 2023.

R Core Team: R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, 2025.

Rauhala, A., Meriö, L.-J., Kuzmin, A., Korpelainen, P., Ala-aho, P., Kumpula, T., Kløve, B., and Marttila, H.: Measuring the spatiotemporal variability in snow depth in subarctic environments using UASs – Part 1: Measurements, processing, and accuracy assessment, The Cryosphere, 17, 4343–4362, https://doi.org/10.5194/tc-17-4343-2023, 2023.

Richiardi, C., Siniscalco, C., and Adamo, M.: Comparison of Three Different Random Forest Approaches to Retrieve Daily High-Resolution Snow Cover Maps from MODIS and Sentinel-2 in a Mountain Area, Gran Paradiso National Park (NW Alps), Remote Sens., 15, 343, https://doi.org/10.3390/rs15020343, 2023.

Riggs, G. A., Hall, D. K., and Salomonson, V. V.: A snow index for the Landsat Thematic Mapper and Moderate Resolution Imaging Spectroradiometer, in: Proceedings of IGARSS '94 - 1994 IEEE International Geoscience and Remote Sensing Symposium, IGARSS '94 - 1994 IEEE International Geoscience and Remote Sensing Symposium, 1942–1944 vol.4, https://doi.org/10.1109/IGARSS.1994.399618, 1994.

Rissanen, T., Aalto, A., Kainulainen, H., Kauppi, O., Niittynen, P., Soininen, J., and Luoto, M.: Local snow and fluvial conditions drive taxonomic, functional and phylogenetic plant diversity in tundra, Oikos, 2023, e09998, https://doi.org/10.1111/oik.09998, 2023.

Rixen, C., Høye, T. T., Macek, P., Aerts, R., Alatalo, J. M., Anderson, J. T., Arnold, P. A., Barrio, I. C., Bjerke, J. W., Björkman, M. P., Blok, D., Blume-Werry, G., Boike, J., Bokhorst, S., Carbognani, M., Christiansen, C. T., Convey, P., Cooper, E. J., Cornelissen, J. H. C., Coulson, S. J., Dorrepaal, E., Elberling, B., Elmendorf, S. C., Elphinstone, C., Forte, T. G. W., Frei, E. R., Geange, S. R., Gehrmann, F., Gibson, C., Grogan, P., Halbritter, A. H., Harte, J., Henry, G. H. R., Inouye, D. W., Irwin, R. E., Jespersen, G., Jónsdóttir, I. S., Jung, J. Y., Klinges, D. H., Kudo, G., Lämsä, J., Lee, H., Lembrechts, J. J., Lett, S., Lynn, J. S., Mann, H. M. R., Mastepanov, M., Morse, J., Myers-Smith, I. H., Olofsson, J., Paavola, R., Petraglia, A., Phoenix, G. K., Semenchuk, P., Siewert, M. B., Slatyer, R., Spasojevic, M. J., Suding, K., Sullivan, P., Thompson, K. L., Väisänen, M., Vandvik, V., Venn, S., Walz, J., Way, R., Welker, J. M., Wipf, S., and Zong, S.: Winters are changing: snow effects on Arctic and alpine tundra ecosystems, Arct. Sci., 8, 572–608, https://doi.org/10.1139/as-2020-0058, 2022.

Rouse, J. W., Haas, R. H., Schell, J. A., and Deering, D. W.: Monitoring vegetation systems in the Great Plains with ERTS, NASA Spec Publ, 351.1, 1974.

Rumpf, S. B., Gravey, M., Brönnimann, O., Luoto, M., Cianfrani, C., Mariethoz, G., and Guisan, A.: From white to green: Snow cover loss and increased vegetation productivity in the European Alps, Science, 376, 1119–1122, https://doi.org/10.1126/science.abn6697, 2022.

Shivanna, K. R.: Climate change and its impact on biodiversity and human welfare, Proc. Indian Natl. Sci. Acad., 88, 160–171, https://doi.org/10.1007/s43538-022-00073-6, 2022.

Simoes, R., Souza, F., Zaglia, M., Queiroz, G. R., Santos, R., and Ferreira, K.: Rstac: An R Package to Access Spatiotemporal Asset Catalog Satellite Imagery, in: 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, 7674–7677, https://doi.org/10.1109/IGARSS47720.2021.9553518, 2021.

Stillinger, T., Roberts, D. A., Collar, N. M., and Dozier, J.: Cloud Masking for Landsat 8 and MODIS Terra Over Snow-Covered Terrain: Error Analysis and Spectral Similarity Between Snow and Cloud, Water Resour. Res., 55, 6169–6184, https://doi.org/10.1029/2019WR024932, 2019.

Tadono, T., Ishida, H., Oda, F., Naito, S., Minakawa, K., and Iwamoto, H.: Precise Global DEM Generation by ALOS PRISM, ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci., II–4, 71–76, https://doi.org/10.5194/isprsannals-II-4-71-2014, 2014.

Verpoorter, C., Kutser, T., Seekell, D. A., and Tranvik, L. J.: A global inventory of lakes based on high-resolution satellite imagery, Geophys. Res. Lett., 41, 6396–6402, https://doi.org/10.1002/2014GL060641, 2014.

Viviroli, D., Dürr, H. H., Messerli, B., Meybeck, M., and Weingartner, R.: Mountains of the world, water towers for humanity: Typology, mapping, and global significance, Water Resour. Res., 43, https://doi.org/10.1029/2006WR005653, 2007.

Wood, S. N.: Generalized Additive Models: An Introduction with R, Second Edition, 2nd ed., Chapman and Hall/CRC, New York, 496 pp., https://doi.org/10.1201/9781315370279, 2017.

Wulder, M. A., Masek, J. G., Cohen, W. B., Loveland, T. R., and Woodcock, C. E.: Opening the archive: How free data has enabled the science and monitoring promise of Landsat, Remote Sens. Environ., 122, 2–10, https://doi.org/10.1016/j.rse.2012.01.010, 2012.

Zanaga, D., Van De Kerchove, R., De Keersmaecker, W., Souverijns, N., Brockmann, C., Quast, R., Wevers, J., Grosu, A., Paccini, A., Vergnaud, S., Cartus, O., Santoro, M., Fritz, S., Georgieva, I., Lesiv, M., Carter, S., Herold, M., Li, L., Tsendbazar, N.-E., Ramoino, F., and Arino, O.: ESA WorldCover 10 m 2020 v100, https://doi.org/10.5281/zenodo.5571936, 2021.

Zhang, T. J.: Influence of the seasonal snow cover on the ground thermal regime: An overview, Rev. Geophys., 43, RG4002, https://doi.org/10.1029/2004rg000157, 2005.

Zhu, Z., Wang, S. X., and Woodcock, C. E.: Improvement and expansion of the Fmask algorithm: cloud, cloud shadow, and snow detection for Landsats 4-7, 8, and Sentinel 2 images, Remote Sens. Environ., 159, 269–277, https://doi.org/10.1016/j.rse.2014.12.014, 2015.

Supplementary materials

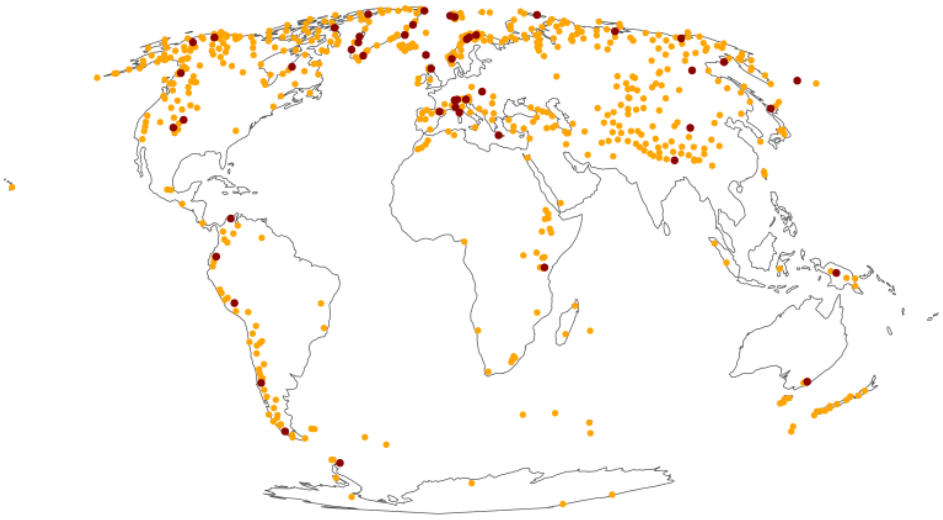

Figure S1. Initial 600 study areas (points in orange) from which the 49 areas (in dark red) were selected for the collection of training data used to fit the Random Forest classifier models.

Supplementary Text S1.

Formulas of the used spectral indices:

*BITM = (((B**2.0) +(G**2.0) +(R**2.0))/3.0) **0.5*

*KBRI = (S1 - N) / (20 * sqrt ((S1 + N)))*

NDSI = (G - S1) / (G + S1)

NDVI = (N - R) / (N + R)

SWM = (B + G) / (N + S1)

where B = blue band, G = green, R = red, N = near infrared, S1 = shortwave infrared 1.

Supplementary Text S2: Usage example

Below, I demonstrate the use of the *snowman* package over Mount Olympus in Greece for the period 2020-2024. The following code introduces all the main functions of the package and produces maps of snow conditions (such as snow cover duration) as output. Running this example with the Lenovo laptop (specified above) with 4 computing cores took 32 minutes (downloads about 60 MB of data).

Before you start, ensure you have the *snowman* package installed. You can install it from GitHub using the following command and load the necessary R packages to run the example:

```
devtools::install_github("poniitty/snowman")
library(terra)
library(snowman)
```

The user is also encouraged to use the most recent version of R and the core dependent packages (such as *terra* and *sf*) because using older versions has resulted in problems with the internal parallelisation of the *snowman* workflow.

Start by setting the number of computer nodes to be utilised in the *snowman* process. The optimal number depends on your machine and the size of the area of interest (AOI). Raster processing in R can be memory-intensive, so processing a large area over a lengthy period with many computer nodes can quickly drain your computer's memory. In this example, I use four cores:

```
n_workers <- 4
```

The user needs to set a directory where to store all the downloaded and processed data and the output maps:

```
landsat_path <- "C:/MyTemp/RS/" # Replace with your own path
```

The algorithm requires an identifier (e.g., name of the study area) and a list of the latitude and longitude of the centre of the AOI in decimal degrees and specifies the size of the buffer around the centre in kilometres that will create the full study area. Here, I use a buffer size of 2 kilometres, which will result in a 4 km x 4 km AOI. Alternatively, the user can provide a simple feature polygon that represents the AOI, in which case the buffer size will be ignored:

```
site <- "Olympus"
aoi_point <- list(lon = 22.349975, lat = 40.083285)
buffer_size <- 2
```

Finally, set the start and end dates for your study period:

```
starting_date <- "2020-01-01"
ending_date <- "2024-12-31"
```

After the presets, run the *extract_landsat* function, which extracts Landsat imagery for the specified AOI and period from the Microsoft Planetary Computer using the rstac package. The output of the extract_landsat function is a data frame with metadata on all downloaded Landsat imagery, which is also used as input in the following functions. Ensure your internet connection is working before running the function. The *snowman_cleanup* will clean the parallelisation session that the future R packages set up for speeding up the downloading process.

```
image_df <- extract_landsat(aoi = aoi_point,
                            site_name = site,
                            aoi_size = buffer_size,
                            start_date = starting_date,
                            end_date = ending_date,
                            sats = c("LT04", "LT05", "LE07", "LC08", "LC09"),
                            base_landsat_dir = landsat_path,
                            workers = n_workers)
snowman_cleanup()
```

Next, the *calc_predictors* function is used to calculate the predictors to be utilised in classifying the Landsat imagery with the Random Forest classifier model:

```
calc_predictors(image_df,
                site_name = site,
                base_landsat_dir = landsat_path)
```

When using the *snowman* package for the first time, or if wanting to ensure that you have the latest versions of the pre-trained classifier models, you can run the *download_model* function, which downloads the classifier model objects from a GitHub repository (https://github.com/poniitty/*snowman*_models):

```
download_model(model_names = c("TM04", "TM05", "LE07", "LC08", "LC09"),
               model_dir = model_path)
```

Once the model objects are downloaded, the classify_landsat function is used to classify the Landsat imagery using the downloaded models into snowy, snow-free, and cloudy pixels:

```
classify_landsat(image_df,
```

```
                site_name = site,
                base_landsat_dir = landsat_path,
                model_dir = model_path,
                workers = n_workers)
```

Finally, the *calc_snow_variables* function calculates the snow variables across the AOI, which are the final products of the *snowman* algorithm. The output is a raster stack including metadata about the Landsat data and the snow variables:

```
snow_vars <- calc_snow_variables(image_df,
                                 site_name = site,
                                 base_landsat_dir = landsat_path,
                                 workers = n_workers)
```

Lastly, the output can be visualised and saved to disc as a GeoTIFF file:

```
plot(snow_vars$scd, col = rev(topo.colors(100)), main = "Snow cover duration in Mount Olympus")
writeRaster(snow_vars, paste0(landsat_path, "/", site, "/", "snow_variables.tif"), datatype = "FLT4S")
```

If the user wants to map the snow over a larger area, my recommendation is to split the whole AOI into subpolygons (e.g., with the *st_make_grid* function from the *sf* R package) and run the *snowman* over the subpolygons one at a time and combine the output at the end of the full process chain.