



HAL
open science

Weakly supervised learning for snow cover segmentation in mountainous areas from Sentinel-1 SAR images using interpolated NDSI time series

Swann Briand, Flora Weissgerber, Sylvain Lobry, Jérôme Idier

► To cite this version:

Swann Briand, Flora Weissgerber, Sylvain Lobry, Jérôme Idier. Weakly supervised learning for snow cover segmentation in mountainous areas from Sentinel-1 SAR images using interpolated NDSI time series. 2025. <hal-05186268>

HAL Id: hal-05186268

<https://hal.science/hal-05186268v1>

Preprint submitted on 25 Jul 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Weakly supervised learning for snow cover segmentation in mountainous areas from Sentinel-1 SAR images using interpolated NDSI time series.

Swann Briand,^{a,c,*}, Flora Weissgerber^a, Sylvain Lobry^b, Jérôme Idier^c

^a*DTIS, ONERA, Université Paris-Saclay, Palaiseau, France*

^b*Université Paris Cité, LIPADE, F-75006 Paris, France*

^c*Nantes Université, École Centrale Nantes, CNRS, LS2N, Nantes, France*

Abstract

Snow cover plays a fundamental role in climate regulation and hydrological processes. Existing snow products are based on optical imagery. Yet, snow monitoring remains challenging in mountainous regions due to frequent cloud cover. Synthetic Aperture Radar (SAR) imagery, unaffected by clouds, enables regular wet snow observations. However dry snow remains mostly transparent to SAR. In this study, we propose a fully automated framework to transfer the snow detection capabilities of optical images, that does not differentiate dry and wet snow, to SAR images. A convolutional neural network is trained to predict a binary snow cover map from a Sentinel-1 Single Look Complex (SLC) dual-pol amplitude image and a snow-free reference image. We threshold the MODIS Normalized Difference Snow Index (NDSI) product to generate binary training labels. Our model is trained in a weakly supervised manner by filling the cloud-induced gaps via temporal interpolation. We first evaluate the influence of the input SAR channels configuration and show that concatenating the acquisition of the day with the reference image is preferable to more complex preprocessing. Then, we compare the Closest Neighbours Interpolation and the Kalman smoother to fill the cloud-induced gaps in the MODIS NDSI time series. We show that increasing the level of supervision improves the model performance. By removing all the gaps and the noise in the NDSI time series, the Kalman smoother yields the best model performance. However the regularization strength of the Kalman smoother is shown to be critical. To validate our method, we compare it to existing snow products. By comparing with the THEIA L2B Snow product, we show that our method gives comparable results to Sentinel-2 based snow cover maps. The comparison with the Copernicus Wet/Dry Snow product shows that our model can detect both wet and dry snow solely from Sentinel-1

*Corresponding author. Email: swann.briand@onera.fr

dual-pol amplitude images. Overall, this study highlights how the snow detection capabilities of optical sensor can be transferred to SAR images thanks to a deep learning framework, to ensure a robust monitoring of wet and dry snow in cloud-prone alpine regions.

Keywords: Snow cover; Optical; SAR; Time series; Pseudo-labels; Weak supervision; Kalman smoother

1. Introduction

Earth’s cryosphere is one of the main drivers of its climate variability, due to its high albedo (0.8 to 0.9 for fresh snow). Indeed, snow covers up to 49 % of the Northern Hemisphere land surface during midwinter time and it stores 75 % of fresh water on land (Lemke et al., 2007). Mountain ranges are considered as the ”water-towers of the world” (Viviroli et al., 2007), providing freshwater for the adjacent areas downstream during spring and summer. Seasonal meltwater sustains ecosystems, irrigates agriculture, and drives hydroelectric production. In Europe, the Alps are the most important water-supplying mountain range (Immerzeel et al., 2019). The Rhône, Danube, Rhine, and Po basins are its main drainage basins and around 170 millions people reside around them (Mastrotheodoros et al., 2020).

To enable the assessment of climate change down to regional and local scales, the Global Climate Observation System program (Global Climate Observing System, 2006) set a requirement of at least 500 m horizontal resolution and 24h temporal resolution to map snow areal extent. Beyond climate monitoring, other applications require more detailed information about the snowpack, including snow depth and Snow Water Equivalent (SWE). These parameters are essential for managing hydroelectric power generation, irrigation systems, and water resources (Rouhier, 2018).

These variables were originally estimated from ground measurements collected at meteorological stations or during field expeditions. While these methods provide highly accurate data, they are spatially limited. To assess snow dynamics at regional scales, geostatistical approaches have been used to interpolate between these sparse data points (Hosang and Dettwiler (1991), Phillips et al. (1992), Erxleben et al. (2002)).

Remote sensing methods overcome these limitations by providing regular observations over large spatial extents and have been extensively used for snow monitoring. Snow detection from optical data leverages the high reflectance of snow in the visible domain and its low reflectance in the Short-Wave InfraRed (SWIR) range (Dozier, 1984, 1989). Hall et al. (1995) introduced the Normalized Difference Snow Index (NDSI), which uses the reflectance contrast between these

two spectral regions to identify snow-covered pixels based on a threshold value. However, cloud cover often obstructs the surface. Even if the NDSI can separate clouds from snow, information about the ground is missing. To address this limitation, several approaches have been developed to fill cloud-induced gaps, including temporal interpolation (Hall et al., 2021), spatio-temporal interpolation (Li et al. (2020), Deng et al. (2024)), optical sensor fusion (Moreno-Martínez et al. (2020), Sedano et al. (2014)), and deep learning techniques (Tsardanidis et al., 2024; Wang et al., 2025).

Synthetic Aperture Radar (SAR), on the other hand, is an active remote sensing technique that illuminates the surface with electromagnetic waves and measures the backscattered signal. Using wavelengths from 0.1 to 100 cm allows it to penetrate clouds. SAR data has been investigated to map snow since the launch of ERS-1 in 1991. The type of information that can be retrieved from SAR data depends on both the characteristics of the data and the detection method used. Wet snow, which strongly attenuates backscattered SAR signal amplitude due to its liquid water content, can be mapped with a dual-polarisation image by comparing the signal attenuation from wet snow with a snow-free reference image (Nagler and Rott, 2000; Nagler et al., 2016), and thresholding the attenuation value. Dynamical thresholding techniques were later introduced to improve classification (James et al., 2024). Monitoring the changes in coherence between two acquisitions, InSAR-based methods can be used to assess the apparition of both dry and wet snow, even in vegetated terrain (Strozzi et al., 1999; Singh et al., 2008). PolSAR techniques that use all four polarization channels (VV, VH, HH and HV) enables to map both dry and wet snow cover (He et al., 2017). Tsai et al. (2019) mapped snow cover using dual-pol images under the assumption of reciprocal targets (VH and HV are considered the same). Despite dry snow being almost transparent to dual-polarization SAR signal, Lievens et al. (2019) proposed a method to retrieve dry snow depth with Sentinel-1 signal depolarization, though it requires prior information on snow presence which is obtained from Sentinel-2 optical data.

To leved the threshold fixing problem, Gallet et al. (2024) applied machine learning techniques to detect wet snow with labels from snowpack model CROCUS (Brun et al., 1989). Lê et al. (2023) predicts wet snow cover map from Sentinel-1 backscatter amplitude and reference image with a Convolutional Neural Network (CNN), using labels derived from Nagler’s method and optical data. Montginoux et al. (2023) predict total (dry and wet) snow cover maps from Sentinel-1 backscatter amplitude using a Convolutional Neural Network (CNN). Its approach builds on the amplitude-based methods of Nagler and Rott (2000) and Lievens et al. (2019),

extending them to total snow cover estimation. This approach relies solely on SAR data as input, and supervision during model training is based on snow cover labels derived from MODIS optical data. Daudt et al. (2023) uses both Sentinel-1 backscatter amplitude and Sentinel-2 optical data, along with elevation information, as input to a Recurrent Neural Network (RNN) to retrieve snow depth at weekly temporal resolution.

In this paper, we present a method to detect both dry and wet snow from SAR dual-pol amplitude images only, using a CNN. Labels are obtained from optical images, allowing us to detect both wet and dry snow without manual labeling. However, cloud cover limits optical labels availability for model training. We propose to use temporal interpolation to create pseudo-labels and increase the level of supervision. To this aim:

- we present a dataset containing SLC coregistered images over different orbits with total snow cover labels
- we investigate whether a weakly supervised method relying on temporal interpolation to increase the amount of optical-derived labels can improve the prediction of total snow cover maps from SAR backscatter amplitude image
- we show that both dry and wet snow can be detected with similar performance from SAR dual-pol amplitude images using a deep-learning framework

The CNN training framework is illustrated Figure 1. In Section 2, we present the temporal interpolation methods and our weakly supervised deep learning framework. In Section 3, we present our study area and the data available, how we process SAR data and create our dataset. In Section 4, we present the model training and the different experiments. In Section 5 we present the results of each experiment and we discuss them in Section 6. Finally, we conclude in Section 7.

2. Methodology

Let $\mathbf{O} \in \mathbb{R}^{T_{\mathbf{O}} \times H_{\mathbf{O}} \times W_{\mathbf{O}}}$ be the NDSI time series and $\mathbf{S} \in \mathbb{R}^{T_{\mathbf{S}} \times H_{\mathbf{S}} \times W_{\mathbf{S}} \times 2}$ be the SAR time series, over the same geographical area. The SAR series \mathbf{S} contains co- and cross-polarized backscatter amplitudes (σ_{VV}^0 and σ_{VH}^0 respectively) in linear scale. Since NDSI and SAR data may differ in temporal and spatial resolutions, we denote $T_{\mathbf{O}}$ and $T_{\mathbf{S}}$ as the number of time steps in the NDSI and SAR time series, and $H_{\mathbf{O}}, W_{\mathbf{O}}$ and $H_{\mathbf{S}}, W_{\mathbf{S}}$ as their respective spatial dimensions (height and width). An index in the time series arrays \mathbf{O} or \mathbf{S} corresponds to a

specific dimension (time, height, or width). Fixing one or more indices defines a position in the array and selects the corresponding value or subset. By fixing the temporal index t , we obtain a full spatial image at date t : $\mathbf{O}_t \in \mathbb{R}^{H_O \times W_O}$ for NDSI data, and $\mathbf{S}_t \in \mathbb{R}^{H_S \times W_S \times 2}$ for SAR data. By fixing the spatial indices (h, w) , we obtain a temporal profile for that pixel: $\mathbf{o}_{h,w} \in \mathbb{R}^{T_O}$ for NDSI, and $\mathbf{s}_{h,w} \in \mathbb{R}^{T_S \times 2}$ for SAR. Fixing all three indices (t, h, w) gives a single observation: $o_{t,h,w} \in [-1, 1]$ for NDSI, and $s_{t,h,w} \in \mathbb{R}^2$ for SAR.

Our method aims to train a Convolutional Neural Network (CNN) f to predict binary snow cover maps $\hat{\mathbf{Y}}_t \in \mathbb{B} = \{0, 1\}^{H_S \times W_S}$ from SAR images $\mathbf{S}_t \in \mathbb{R}^{H_S \times W_S \times 2}$. To supervise our network training for this binary segmentation problem, we derive labels from the NDSI time series. More precisely, we use a threshold for the NDSI values to obtain binary snow cover labels. More details about the NDSI are given in Section 2.1. However, due to cloud cover, some pixels $o_{h,w,t}$ cannot be observed. We propose to reconstruct the NDSI time series \mathbf{O} by filling gaps in each pixel time series $\mathbf{o}_{h,w} = (o_0, \dots, o_{T_O-1})_{h,w}$ using temporal interpolation. Two solutions are explored: a simple Closest Neighbours Interpolation (CNI) method, and a more advanced one based on Kalman smoothing (KS), which fills all the gaps.

Although both SAR and NDSI data are structured as time series, our CNN is applied independently for each date t , without modeling temporal dependencies. The SAR time series is used solely to build the multi-channel input \mathbf{X}_t . The NDSI time series serves only for supervision during training to generate labels, or pseudo-labels via interpolation.

In this section, we first introduce the NDSI and its uses. We then present the two temporal interpolation methods used to generate pseudo-labels, and finally we describe our model training framework.

2.1. Normalized Difference Snow Index

Snow is highly reflective in the visible part of the electromagnetic spectrum and highly absorptive in the short-wave infrared domain, whereas cloud reflectance remains high in both spectral regions. Based on these spectral properties, the Normalized Difference Snow Index (NDSI) was introduced to distinguish snow from clouds in optical imagery by combining reflectance values from the visible (typically green) and SWIR bands (Dozier (1984), Hall et al. (1995)).

$$\text{NDSI} = \frac{\rho_{\text{Green}} - \rho_{\text{SWIR}}}{\rho_{\text{Green}} + \rho_{\text{SWIR}}} \quad (1)$$

where ρ denotes surface reflectance. NDSI ranges between -1 and 1, the presence of snow in a pixel theoretically lead to a positive NDSI. However Hall et al. (1995) introduced a 0.4 threshold

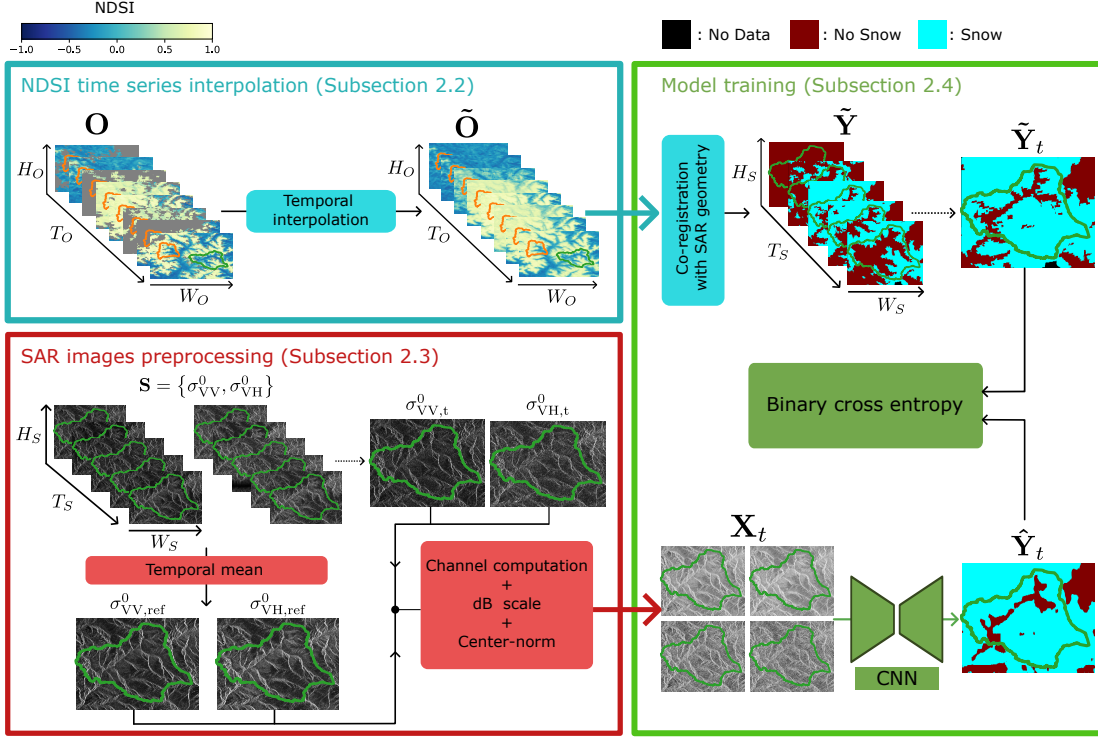


Figure 1: Illustration of our training framework. Dotted arrows show the selection of one date among the time series. The optical NDSI time series \mathbf{O} is temporally interpolated to obtain a cloud-gap filled time series $\tilde{\mathbf{O}}$. The interpolation methods are described in Subsection 2.2. The time series $\tilde{\mathbf{Y}}$ consists of binary snow cover pseudo-labels. It is obtained by thresholding the interpolated NDSI time series $\tilde{\mathbf{O}}$ and aligning it spatially and temporally onto the SAR time series \mathbf{S} . \mathbf{S} contains backscatter amplitude values in VV and VH polarisations (σ_{VV}^0 and σ_{VH}^0 respectively). To create the SAR inputs, we first compute reference images for both polarisations $\sigma_{VV,\text{ref}}^0$ and $\sigma_{VH,\text{ref}}^0$. We estimate snow cover maps at time t by selecting σ_{VV}^0 and σ_{VH}^0 values from this date, $\sigma_{VV,t}^0$ and $\sigma_{VH,t}^0$, combine it to the reference images to build the input channels \mathbf{X}_t . The SAR data preprocessing and different channels combinations are described in Subsection 2.3. To train the model f , we minimize a binary cross entropy between the estimated snow cover map at time t $\tilde{\mathbf{Y}}_t$ and model prediction $\hat{\mathbf{Y}}_t = f(\mathbf{X}_t)$. The NDSI data is used only during training. The training framework is described in Subsection 2.4. For inference, only SAR data is needed to estimate snow cover maps.

to assess if a pixel is snow covered or snow free. They empirically determined that if a pixel has a NDSI value above 0.4, more than half of its area is covered with snow. Several studies optimize the threshold which can be suboptimal for pixel in dark or forested areas or if it covers a too large area (Härer et al., 2018). Hall and Riggs (2021) uses additional data screens like altitude, surface temperature and other reflectance bands to filter out pixels that have a positive NDSI but no snow cover. The NDSI can be regressed to retrieve finer-grained information, such as fractional snow cover, the snow-covered fraction of a pixel area (Salomonson and Appel, 2004; Gascoïn et al., 2020), Snow Water Equivalent or snow depth (Teleubay et al., 2023). In our method, we apply the traditional 0.4 threshold to derive binary snow cover labels.

2.2. Optical time series interpolation

2.2.1. Closest Neighbours Interpolation

As a simple temporal interpolation method, we perform linear interpolation to fill gaps in a three-days window, which is the mean duration of cloud obstruction (Hou et al., 2019).

$$\begin{cases} \tilde{o}_t = \frac{1}{2}(o_{t-1} + o_{t+1}) \\ \tilde{o}_t = o_{t-1} \text{ if } o_{t+1} \text{ missing} \\ \tilde{o}_t = o_{t+1} \text{ if } o_{t-1} \text{ missing} \\ \text{NaN else} \end{cases} \quad (2)$$

In the remaining of this paper, we will refer to this interpolation method as Closest Neighbours Interpolation (CNI). This method is close to the Cloud Gap Filled method used in the MOD10A1F dataset (Hall et al., 2021), which extends the last NDSI measurement until a new observation is available. With CNI, future observations are also taken into account and we minimize the risk of missclassification by extending the measurement only to the next day. However, CNI will not fill all the gaps in the optical time series.

2.2.2. Kalman smoothing

The Kalman smoother (KS) is a recursive algorithm that estimates the successive states of a system given a linear state-space model (SSM) that describes the system dynamics and a serie of observations, by minimising a mean squared-error between predictions and observations (Kalman, 1960). Both model noise and observation noise are assumed to follow a normal distribution of mean zero and known variance. The smoothing implies that information from both past and future observations are used to estimate the state of a pixel. A filtering phase is first applied to use past information, then a backward pass adds future observations information. Kalman filtering and smoothing considers that the state of the system is hidden and that information about the system is only available through noisy observations.

The relationship between hidden states and observation must be linear, as well as the hidden state dynamics between two timesteps. In our case, we obtain our binary labels from NDSI, so both the hidden state and the observations are NDSI. We denote the "true" hidden NDSI time series of a pixel at location (h, w) , $\mathbf{x}_{h,w} = (x_0, \dots, x_{T_{\mathbf{O}}})$ and the associated series of observations $\mathbf{o}_{h,w} = (o_0, \dots, z_{T_{\mathbf{O}}})$. Moreover, snow (and thus NDSI) dynamics are complex and depend on other variables such as wind speed, air temperature or slope. As we expect the NDSI value to be close between two consecutive dates, we model the NDSI dynamic as a random walk with a

step sampled from a gaussian distribution with known mean and variance. The SSM is expressed as:

$$\begin{cases} x_{t+1} &= x_t + \epsilon_t^{\text{mod}} \\ o_t &= x_t + \epsilon_t^{\text{mes}} \end{cases} \quad (3)$$

ϵ_t^{mod} and ϵ_t^{mes} are the model and observation noise. Both are assumed to follow Gaussian distributions with zero mean and known variance, v_{mod} and v_{obs} respectively. We consider the model noise variance v_{mod} as constant, and use it to normalize the estimated variance. We define η the normalized observation noise variance and \bar{v} the normalized estimated noise variance:

$$\eta = v_{\text{obs}}/v_{\text{mod}} \quad (4)$$

$$\bar{v} = v/v_{\text{mod}} \quad (5)$$

To interpolate the time series, the KS first predicts the next state using the SSM and all past information. It then corrects the prediction by integrating observations. Once the forward pass is done, a backward pass is performed to include future observations.

Prediction Step:

$$\hat{x}_{t|t-1} = \hat{x}_{t-1|t-1} \quad (6)$$

$$\bar{v}_{t|t-1} = \bar{v}_{t-1|t-1} + 1 \quad (7)$$

$$\bar{s}_t = \bar{v}_{t|t-1} + \eta \quad (8)$$

Correction Step:

$$\bar{k}_t = \bar{v}_{t|t-1} \bar{s}_t^{-1} \quad (9)$$

$$\hat{x}_{t|t} = \hat{x}_{t|t-1} + \bar{k}_t v_{\text{mod}} (o_t - \hat{x}_{t|t-1}) \quad (10)$$

$$\bar{v}_{t|t} = \bar{v}_{t|t-1} - \bar{k}_t^2 \bar{s}_t \quad (11)$$

Smoothing Step:

$$c_t = \frac{\bar{v}_{t|t}}{\bar{v}_{t|t+1}} \quad (12)$$

$$\hat{x}_{t|T} = \hat{x}_{t|t} + c_t (\hat{x}_{t+1|T} - \hat{x}_{t|t}) \quad (13)$$

$$\bar{v}_{t|T} = \bar{v}_{t|t} + c_t^2 (\bar{v}_{t+1|T} - \bar{v}_{t|t}) \quad (14)$$

When observations are not available for the time step t , the correction step is skipped and the *a priori* estimate becomes the *a posteriori* estimate used for next step prediction. The

estimated noise variance $\bar{v}_{t|t}$ keeps increasing the more timesteps without observation the time series has (Eq. 7). At each timestep, the estimated mean $\hat{x}_{t|t}$ only depends on η and $\bar{v}_{t|t}$, which is scaled by v_{mod} . The variable η acts as a regularization parameter, which is used to compute the Kalman gain k that weights model prediction and observations. We fix the model noise variance to 1% of the max NDSI (i.e. $v_{\text{mod}} = 0.01$) and generate time series for several values of $\eta \in [10^{-3}; 10^2]$. We initialize the smoother with mean $x_0 = o_0$ or 0 if $o_0 = \text{NaN}$ and variance $v_0 = 0.01$. Each pixel at position h, w in an image taken at time t , has an associated mean $\hat{x}_{h,w,t}$ and variance $v_{h,w,t}$. We only consider the estimated mean value and denote the resulting interpolated time series: $\tilde{\mathbf{O}} \in \mathbb{R}^{T_{\mathbf{O}} \times H_{\mathbf{O}} \times W_{\mathbf{O}}}$.

2.3. SAR time series preprocessing

To produce the SAR inputs \mathbf{X}_t used for snow cover prediction, we preprocess the SAR time series \mathbf{S} to derive snow-informative features. We compute reference images $\sigma_{\text{VH,ref}}^0$ and $\sigma_{\text{VH,ref}}^0$ from snow-free observations. This approach to improve SAR-based snow detection was introduced by Nagler and Rott (2000), and has become standard in SAR-based snow studies for detecting wet snow by comparing snow-covered acquisitions to snow-free reference scenes (Lê et al., 2023; James et al., 2024; Gallet et al., 2024; Karbou et al., 2021; Montginoux et al., 2023). Lievens et al. (2019) also used reference images for the estimation of dry snow height. Reference images are obtained by averaging the backscatter amplitude over snow-free summer dates (from July to August). The averaging is performed using a maximum likelihood estimator under a gamma distribution assumption. Before computing the reference, outlier values of \mathbf{S} are saturated at the 0.5th and 99.5th percentiles to reduce speckle noise influence and \mathbf{S} is converted to dB scale.

These reference images allow us to build channels which are less sensitive to local acquisition geometry. Different channel combinations such as raw backscatter amplitude, backscatter-to-reference ratios or concatenation, or snow-specific indices like R_{wet} and R_{dry} (Eqs. 19–20) are compared in terms of their effectiveness for snow detection. The construction and selection of these SAR channels are detailed in Subsection 4.2.

2.4. Model training framework

We train a Convolutional Neural Network (CNN) f to predict binary snow cover maps $\hat{\mathbf{Y}}_t \in \mathbb{B}^{H_s \times W_s}$ for a given date t , using as input a single SAR image $\mathbf{X}_t \in \mathbb{R}^{H_s \times W_s \times C}$ derived from the SAR time series. Although the full SAR time series \mathbf{S} is used during preprocessing to construct the input \mathbf{X}_t to compute temporal reference values, the model operates independently

on each date, the full time series is never input to the model.

$$\hat{\mathbf{Y}}_t = f(\mathbf{X}_t) \tag{15}$$

The model is supervised using pseudo-labels derived from the interpolated optical NDSI time series $\tilde{\mathbf{O}}$. We first perform temporal alignment by retaining only the dates t where a SAR observation \mathbf{S}_t is available. For each of these dates, we project the corresponding NDSI map $\tilde{\mathbf{O}}_t$ onto the SAR geometry using georeferencing pixel information obtained by the projection of the RGE alti 1m (Institut national de l’information géographique et forestière (IGN), 2021) in the French part of the image and the TINITALY 10 m (Tarquini et al., 2007) in the Italian part of the image. When projecting the RGE alti 1m, multiple DEM points are projected in the same pixel. The georeferencing information are obtained by averaging these points. On the Italian part of the image, the projection of the DEM leads to a sparse map that is then interpolated. The aligned NDSI maps are then thresholded at 0.4 to generate a binary snow cover time series $\tilde{\mathbf{Y}} \in \mathbb{B}^{T_s \times H_s \times W_s}$.

During training, the model receives a SAR input \mathbf{X}_t for a given date t and predicts a binary snow cover map $\hat{\mathbf{Y}}_t$. The model is trained by minimizing the binary cross-entropy loss between prediction and target map $\tilde{\mathbf{Y}}_t$:

$$\mathcal{L}_{BCE}(\hat{\mathbf{Y}}_t, \tilde{\mathbf{Y}}_t) = - \sum_{r=1}^{H_s \times W_s} [\tilde{y}_r \log(\hat{y}_r) + (1 - \tilde{y}_r) \log(1 - \hat{y}_r)] \tag{16}$$

We ignore pixels without labels when summing the loss. Note that the optical data is used only during training to provide supervision. Once trained, the model predicts snow cover maps only from SAR data.

3. Data and study area

3.1. Study area

We study the Guil and Gyronde river basins located in the French Alps near the Italian border (illustrated in Figure 2) during the 2018 and 2019 study years, defined here as running from July 1st, 2018 to June 30th, 2019, and from July 1st, 2019 to June 30th, 2020, respectively. Each period includes the summer preceding the water year (which spans from September 1st to August 31st), allowing us to capture pre-season conditions relevant for snow and soil dynamics. The Guil basin covers an area of 411 km², while the Gyronde basin covers an area of 230 km². Figure 3 illustrates the elevation profiles of both basins. The Guil basin has an altitude ranging homogeneously from 1000 m to 4000 m while Gyronde basin exhibits steeper slope. In

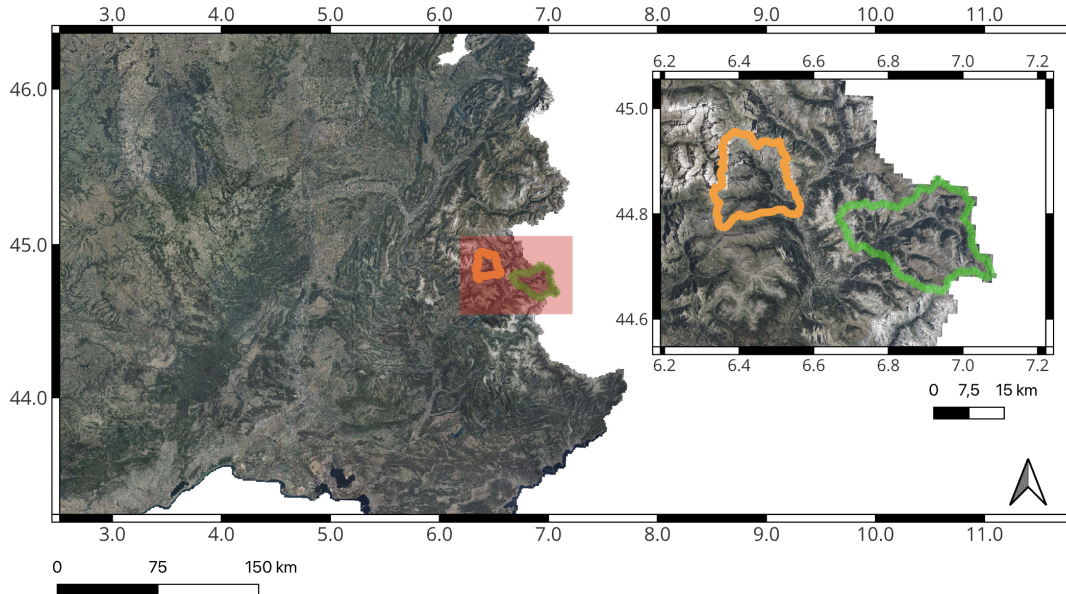


Figure 2: Map of the two study basins in Rhône-Alpes region, France. —: Guil basin, —: Gyronde basin.

both basins, elevations are predominantly concentrated around 2500 meters. This mid-altitude configuration leads in complex snow dynamics. At these elevations, temperatures frequently fluctuate around the freezing point, leading to multiple freeze-thaw cycles. These cycles influence the snowpack stability, soil moisture and root growth Man et al. (2022). Consequently, our remote sensing method needs to be able to map both wet and dry snow to allow precise alpine environmental dynamics studies.

In this study, we use data from the Guil basin for both 2018 and 2019 periods, and from the Gyronde basin for 2018 only. The Guil 2018 time series is used as training and validation sets, and test set without domain shift. The Guil 2019 and Gyronde 2018 are used as two other test sets to test our method robustness to temporal and spatial domain shift.

3.2. Sentinel-1 SAR images

We use Sentinel-1 SLC IW (Single Look Complex - Interferometric Wide swath) data acquired by the Sentinel-1A and Sentinel-1B satellites within the European Space Agency’s Copernicus program. The acquisitions we use are performed in C-band (~ 5.4 GHz), and provide dual-polarization (VV and VH) acquisitions with a spatial resolution of approximately 5×20 m in range-azimuth. The temporal revisit frequency of 6 days allows for regular monitoring of snow cover dynamics. The images are debursting and coregistered with an interferometric precision using LabSAR algorithm (Weissgerber et al., 2022). Three Sentinel-1 orbits cover the




Figure 3: Elevation profile of Guil and Gyronde basins.
■: Guil, ■: Gyronde,

Guil basin: two descending orbits D139 and D66, and one ascending orbit A88. The Gyronde basin is also covered by three orbits: descending orbit D139 and ascending orbits A88 and A161. Therefore, we have 3 acquisitions every 6 days for each basin. The spatial dimensions (H_S, W_S) of the Sentinel-1 images vary across acquisition dates, as they depend on the specific orbit and imaging geometry. As a result, each time series is composed of temporally ordered SAR images with date-dependent spatial extents. Table 1 summarizes the Sentinel-1 acquisitions used for each time series.

Basin and year	Orbit	Number of dates	Spatial dimensions (Height \times Width in pixels)
Guil 2018	D139	55	1872 \times 8889
	D66	55	1996 \times 7833
	A88	54	1934 \times 10745
Guil 2019	D139	56	1872 \times 8889
	D66	56	1996 \times 7833
	A88	54	1934 \times 10745
Gyronde 2018	D139	54	2038 \times 4214
	A88	54	1964 \times 5107
	A161	28	2147 \times 6019

Table 1: Sentinel-1 acquisitions used over the Guil and Gyronde basins during the 2018 and 2019 study periods.

3.3. Snow products

3.3.1. MODIS NDSI

The MODIS instrument provides multispectral observations every 1 to 2 days over the globe. We use the MOD10A1 (Hall and Riggs, 2021), which provides daily NDSI measurements over our study area at 500×500 m resolution at nadir. The one day revisit time ensures that we always have labels from the same day as each SAR acquisition. The NDSI product is a composite of several acquisitions of the day. In addition to NDSI, the dataset provides a cloud mask. The binary labels are obtained by thresholding the NDSI at 0.4. For the pseudo-labels, the thresholding is applied after the NDSI time series interpolation.

3.3.2. THEIA L2B Snow product

The THEIA L2B Snow product provides a binary snow cover product and a cloud mask based on Sentinel-2 and Landsat-8 data at 20 m resolution. An acquisition is available every 5 or 10 days. The snow product is based on the Let-it-snow (Simon et al., 2018) processing chain, which takes as input Sentinel-2 and Landsat-8 surface reflectances corrected for atmospheric and slope effects, a digital elevation model and the associated cloud mask from THEIA L2A product. The Sentinel-2 tiles covering our study area are 32TLQ and 31TGK.

THEIA L2B Snow product is only used to validate our method, and to assess the transfer of the total snow detection capacity from optical images to SAR images.

3.3.3. Copernicus Wet/Dry Snow product

The Copernicus Wet/Dry Snow (WDS) product provides snow state classification over Europe in the entire EE38+UK area at 60 m resolution. This dataset is based on overlapping acquisitions from Sentinel-1 and Sentinel-2 at the same date. To detect total snow, it uses the Sentinel-2 Fractional Snow Cover on the Top Of Canopy (FSCTOC). Binary snow cover extent is retrieved from FSCTOC by applying a threshold at 60%. To differentiate wet snow from dry snow, it uses the SAR-based algorithm on Sentinel-1 data to map snowmelt area extent with Nagler’s method (Nagler et al., 2016). A pixel is classified as wet snow when detected with S1 wet snow algorithm and $FSCTOC \geq 60\%$. When $FSCTOC \geq 60\%$ and no wet snow is detected with S1, the pixel is classified as dry snow. Forested areas, urban areas and water areas are masked using the Copernicus High Resolution Layers Tree Cover Density (TCD) of the year 2018, Imperviousness Density (IMD) of the year 2018 and the Water Layer of 2018 because the wet snow detection was deemed unreliable. Clouds, clouds shadows and regions affected by SAR shadows, foreshortening and layover are also masked.

Copernicus WDS products is not used for training our model, but to validate our model ability to detect both wet and dry snow.

4. Experimental settings

4.1. Model training and evaluation

We train the network using Guil basin acquisitions between September 1st, 2018 and June 1st, 2019. To compute reference images, we use acquisitions from July 1st to August 31st, 2018, as detailed in Subsection 4.2. Out of the 134 acquisitions available across the three orbits, we use 88 dates for training, 24 for validation, and 22 for testing. Figure 4 shows the split in train validation and test set of the 2018 water year.

We chose the test dates so that a Sentinel-2 acquisition, and thus a label map from THEIA L2B Snow or Copernicus WDS, is also available. Since THEIA L2B Snow and Copernicus WDS products are available every 5 to 10 days, we regularly sample the water year ensuring that our test set is representative of the whole water year.

We extract 512×512 patches from the 88 training images in the 2018 Guil time serie using a regular grid with overlap to increase the amount of training data. A simple U-Net architecture (Ronneberger et al., 2015) is trained using an EfficientNet-B0 encoder. Interpolated labels are used for training only. For validation, non-overlapping 512×512 patches are extracted from the 24 validation images.

Each model is characterised by the set of SAR channels used as inputs and the interpolation method used to label the SAR pixels. For each configuration, we train 6 models. We consider three hyperparameters: the learning rate, the number of training epochs and the classification threshold. The learning rate is shared across all models of the same configuration (*i.e.*, channels set and interpolation method). To prevent overfitting, we use early stopping: training is stopped if the validation loss does not decrease by more than 0.01 points over 5 consecutive epochs. After training, we choose the classification threshold that maximizes the Overall F1-score on the validation set. The number of training epochs and the classification threshold are specific to each model, even within the same configuration.

Predictions are performed over the entire image domain $H_S \times W_S$ using overlapping 512×512 patches with a stride of 128 pixels. Each pixel is predicted multiple times with varying context, and final predictions are obtained by Gaussian-weighted averaging, giving more importance to central pixels.




Figure 4: Guil 2018 time series split into train, validation, and test sets. Orbits are shown with bar color above the middle line and sets with bar colors below. Test dates also have a Sentinel-2 observation available. Orbits: █: D139, █: A88, █: D66. Sets: █: reference dates, █: train dates, █: validation dates, █: test dates.

4.2. Selection of SAR channels

Sentinel-1 SLC images are preprocessed before being used as input to the CNN. We discard the phase and retain only the amplitude-related data. All preprocessing steps described in Subsection 2.3 are applied independently for each polarisation, orbit, basin and water year. To compute these reference images, we select all the dates between July 1st and August 31st and estimate the reference images $\sigma_{\text{VV,ref}}^0$ and $\sigma_{\text{VH,ref}}^0$ with the maximum likelihood estimator of a Rayleigh law.

We compare four sets of SAR channels described in Tab. 2. The set **A** concatenates the two base polarimetric channels. The set **B** adds reference images. The set **C**, also used in L e et al. (2023) and Gallet et al. (2024) concatenates ratios of backscatter channels with their references.

$$\sigma_{\text{VV,ratio}}^0 = \sigma_{\text{VV}}^0 / \sigma_{\text{VV,ref}}^0 \quad (17)$$

$$\sigma_{\text{VH,ratio}}^0 = \sigma_{\text{VH}}^0 / \sigma_{\text{VH,ref}}^0 \quad (18)$$

Using ratios allows to be independent from local incident angle as we measure a difference from the snow-free response for each pixel. The set **D**, used in Montginoux et al. (2023) concatenates specialized channels for wet and dry snow detection R_{dry} and R_{wet} .

$$R_{\text{wet}} = \frac{1}{2}\sigma_{\text{VV,ratio}}^0 + \frac{1}{2}\sigma_{\text{VH,ratio}}^0 \quad (19)$$

$$R_{\text{dry}} = \frac{\sigma_{\text{VH}}^0 / \sigma_{\text{VV}}^0}{\sigma_{\text{VV,ref}}^0 / \sigma_{\text{VH,ref}}^0} \quad (20)$$

The channel R_{wet} comes from Nagler’s method, with a constant factor 0.5 to weight VV and VH polarization equally. The channel R_{dry} is an adaptation of the one introduced in Lievens et al. (2019). The ratio between VV and VH polarizations allows to assess the depolarization caused by the dry snow pack. These channels are optimized for wet and dry snow detection, we want to assess if the network needs specialized channels or finds better features itself for snow detection. Table 2 summarizes the channels sets.

ID	Concatenated channels
A	$\{\sigma_{VV}^0, \sigma_{VH}^0\}$
B	$\{\sigma_{VV}^0, \sigma_{VH}^0, \sigma_{VV,ref}^0, \sigma_{VH,ref}^0\}$
C	$\{\sigma_{VV,ratio}^0, \sigma_{VH,ratio}^0\}$
D	$\{R_{dry}, R_{wet}\}$

Table 2: Description of the compared set of SAR channels.

For each input channel, we compute the mean and standard deviation over the entire time series to standardize the data.

4.3. Robustness to domain shift: spatial and temporal transfer

Deep learning models performance can decrease when generalizing outside of their training domain. However, training data is limited. Our goal is to produce a model that does not need to be retrained for each study area and each new . The training is done only using the SAR and NDSI time series of the 2018 water year over the Guil basin. To evaluate if our method produces models robust to temporal or spatial transfer we use additional test sets. To test the robustness to temporal transfer, we include the SAR and NDSI time series of the 2019 water year over the Guil basin. As to test the robustness to spatial transfer we include the SAR and NDSI time series of the 2018 water year over the Gyrone basin.

The evaluation on MODIS NDSI, THEIA L2B Snow and Copernicus WDS labels will be performed on each three domains.

4.4. Evaluation metrics

4.4.1. Total snow

For quantitative model evaluation on MODIS NDSI and THEIA L2B snow products, we compute the cumulated confusion matrix on all dates of the test time series (*Snow* is the positive class). Accuracy measures the proportion of correct predictions for both classes. Precision is the amount of correct predictions among all predictions of a class, recall is the amount of correct predictions among all actual instances of a class. The F1-score is the harmonic mean between precision and recall, and the Overall F1-score is the average of F1-score for positive and negative class weighted by the frequency of each class. Because we use the cumulated confusion matrix

on the whole time series our F1-score is close to a micro-averaged F1-score.

$$\begin{aligned}
Precision_{\text{Snow}} &= \frac{TP}{TP + FP} \\
Recall_{\text{Snow}} &= \frac{TP}{TP + FN} \\
Precision_{\text{NoSnow}} &= \frac{TN}{TN + FN} \\
Recall_{\text{NoSnow}} &= \frac{TN}{TN + FP}
\end{aligned} \tag{21}$$

$$\begin{aligned}
Accuracy &= \frac{TP + TN}{TP + FP + TN + FN} \\
\text{F1-score} &= \frac{2 \times Precision \times Recall}{Precision + Recall}
\end{aligned} \tag{22}$$

$$\begin{aligned}
\text{Overall F1} &= f_{\text{Snow}} \times \text{F1-score}_{\text{Snow}} + \\
&\quad (1 - f_{\text{Snow}}) \times \text{F1-score}_{\text{NoSnow}}
\end{aligned}$$

with TP the number of true positives, FP the false positives, TN the true negatives and FN the false negatives. Model scores are given as mean and standard deviation, computed among the 6 models with identical configurations.

4.4.2. Wet and dry snow

As our model predicts binary snow, we cannot compute all the metrics on wet and dry snow when evaluating on Copernicus WDS product. We consider a pixel as true positive when its label is wet or dry snow and the model predicts it as snow covered. We consider a pixel as false negative when its label is wet and dry snow and the model predicts it as snow free.

$$\begin{aligned}
TP_{\text{wet}} &= \sum_i [(\hat{y}_i = \text{Snow}) \wedge (y_i = \text{WetSnow})] \\
FN_{\text{wet}} &= \sum_i [(\hat{y}_i \neq \text{Snow}) \wedge (y_i = \text{WetSnow})] \\
TP_{\text{dry}} &= \sum_i [(\hat{y}_i = \text{Snow}) \wedge (y_i = \text{DrySnow})] \\
FN_{\text{dry}} &= \sum_i [(\hat{y}_i \neq \text{Snow}) \wedge (y_i = \text{DrySnow})]
\end{aligned} \tag{23}$$

We cannot compute false positives and true negatives for dry and wet snow classes. We can thus only evaluate our model recall on the following classes:

$$\begin{aligned}
Recall_{\text{wet}} &= \frac{TP_{\text{wet}}}{TP_{\text{wet}} + FN_{\text{wet}}} \\
Recall_{\text{dry}} &= \frac{TP_{\text{dry}}}{TP_{\text{dry}} + FN_{\text{dry}}}
\end{aligned} \tag{24}$$

5. Results

5.1. Selection of SAR channels

SAR channel set	Accuracy \uparrow	Overall F1 \uparrow	F1 _{Snow} \uparrow	F1 _{No Snow} \uparrow
A: $\{\sigma_{VV}^0, \sigma_{VH}^0\}$	0.872 (0.0043)	0.873 (0.0042)	0.897 (0.0038)	0.831 (0.0054)
B: $\{\sigma_{VV}^0, \sigma_{VH}^0, \sigma_{VV,ref}^0, \sigma_{VH,ref}^0\}$	0.901 (0.0044)	0.902 (0.0042)	0.921 (0.0041)	0.867 (0.0044)
C: $\{\sigma_{VV,ratio}^0, \sigma_{VH,ratio}^0\}$	0.876 (0.0042)	0.876 (0.0039)	0.901 (0.0038)	0.832 (0.0044)
D: $\{R_{dry}, R_{wet}\}$	0.876 (0.0080)	0.877 (0.0077)	0.901 (0.0070)	0.835 (0.0090)

Table 3: Mean performance and associated standard deviations of models trained with different SAR channel sets on the Guil 2018 test set.

Table 3 shows model performance on Guil 2018 test set. The mean and standard deviation are computed by evaluating metrics on 6 models trained with different seeds. This variability helps assess the robustness of the models and avoids overestimating performance due to a favorable seed. Using only the two backscatter amplitude channels σ_{VV}^0 and σ_{VH}^0 (channel set **A**) yields an accuracy above 0.87, an F1-score of 0.897 for the class *Snow* and 0.831 for the class *NoSnow*. These results demonstrate that our method can accurately predict both dry and wet snow cover maps only from SAR data. Adding reference from means of snow-free images significantly improves accuracies and F1 for all models, whether it is concatenated with the raw backscatter amplitude σ_{VV}^0 and σ_{VH}^0 (channel set **B**) or combined with a ratio (**C** and **D**). Channel set **B** improves performance by 3 points on all metrics. Channel sets **C** and **D** yield smaller scores, but still improve from **A**. Sets **A**, **B**, and **C** show similar variance, while set **D** introduces slightly higher variability.

For qualitative evaluation, we look at model predictions on the Guil basin on March 31st 2019 in Figure 5. All channel sets allow models to distinguish the majority of snow-covered from snow-free regions. Channel set **A** contains some speckle noise and the predicted snow cover map misses some *NoSnow* patches. Channel sets **C** and **D** are highly noised, and lead to poor segmentation with snow cover maps showing scattered and noisy predictions, especially in the center of the basin. Whereas channel set **A** yields compact patches, predictions from sets **C** and **D** appear more granular and result in numerous localized misclassifications. This

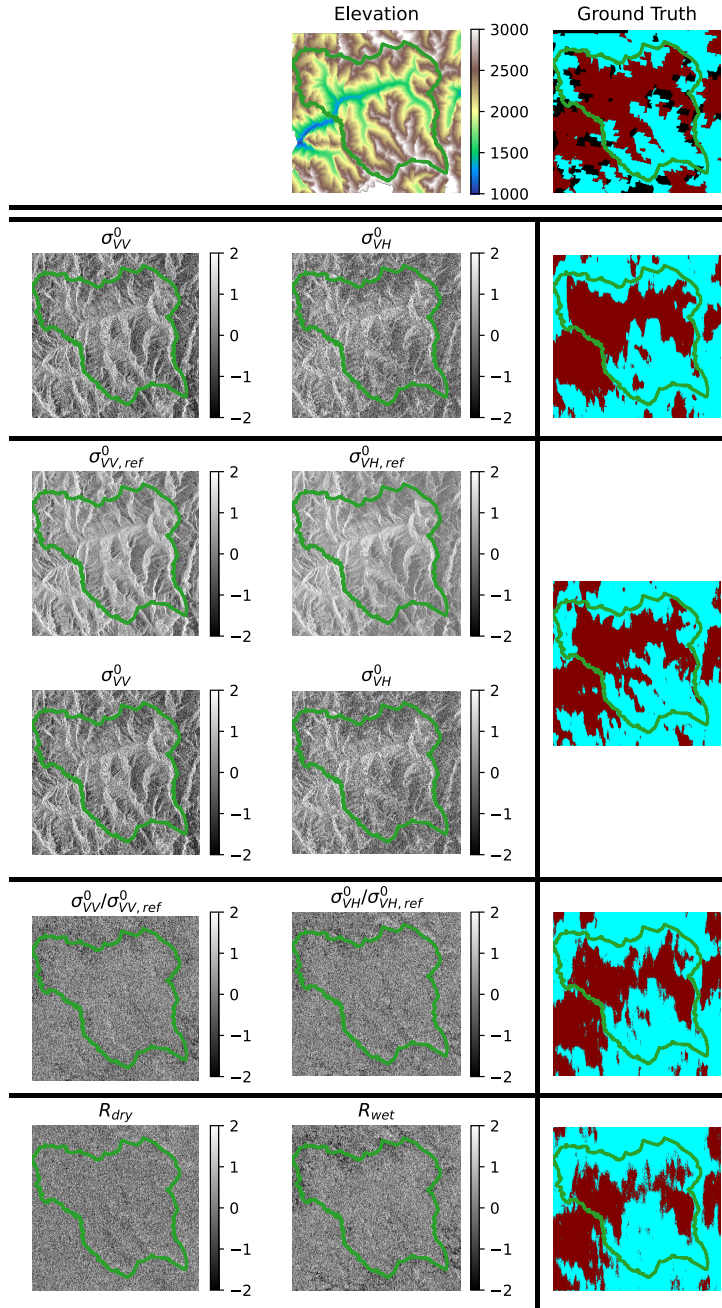


Figure 5: Predictions in the Guil basin from acquisitions on March 31st, 2019, from models trained with different SAR channels. The ground truth is displayed in the top-right corner. SAR channels are centered and normalized in dB, predicted snow maps are in the right column. —: Guil basin, cyan: *Snow*, red: *No Snow*, black: *NoData*.

behavior resembles that of Nagler’s method when a fixed threshold is applied without prior image filtering. The ground truth shows that the spatial distribution of snow cover is correlated to the

terrain elevation as intuition would suggest. **C** and **D** use ratios between backscatter amplitude channels and their references to create slope-independent features, but remove topographical information. While this can help reduce terrain-related variability, it tends to remove useful topographic information and leads to imprecise segmentation. We see that Channel set **B**, which includes reference images without ratios, allows the best segmentation, even if the *NoSnow* class is slightly over-predicted.

Based on both quantitative metrics and qualitative analysis, we select SAR channel set **B** for the remaining experiments.

5.2. Weak supervision through label interpolation

We increase the level of supervision by temporally interpolating the NDSI time series **O** and projecting it on the SAR training set. We compare our two interpolation methods, CNI and KS-based. To select the optimal level of regularization for KS without accessing a ground truth, we treat η as a hyperparameter and train models with various η values. Their performance is then evaluated on a non-interpolated validation set, which is reserved for this purpose.

5.2.1. Impact of the interpolation methods on the snow product time series

Interpolation	None	CNI
NoData (%)	51.81	22.60
Snow (%)	22.26	42.26
No Snow (%)	25.93	34.73
f_{Snow}/f_{NoSnow}	1.16	1.23

Table 4: Class distribution on MODIS geometry in Guil basins for 2018 water year. f_{Snow} is the frequency of the *Snow* class in the dataset, f_{NoSnow} the frequency of the *No Snow* class.

Table 4 shows *Snow* and *NoSnow* class distribution on the Guil basin during the 2018 water year. The CNI method is enough to reduce by half the amount of missing labels. However, it also introduces a shift in the class distribution from observations by increasing the occurrence of the *Snow* class. This shift is likely due to the fact that snowfall events occur with cloud cover. As fresh snow exhibits a high NDSI, due to its low liquid water content and strong reflectance in the visible spectrum, the interpolation tends to yield elevated NDSI values, favoring *Snow* classification.

Figure 6 compares the interpolated NDSI time series of a MODIS pixel obtained using the CNI method and the Kalman Smoother (KS) for a range of η values, along with the corresponding estimated variance. Blue dots represent the NDSI values estimated by the CNI method. As expected from its design, CNI keeps the interpolated values close to the original observations

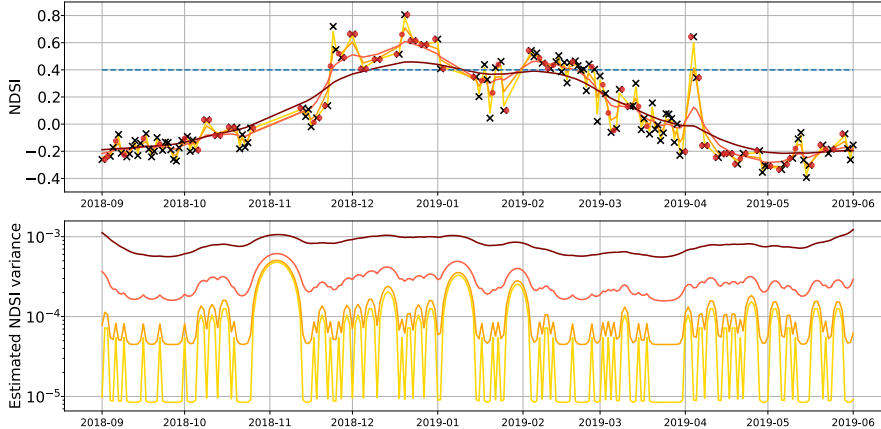


Figure 6: Interpolated NDSI time series of a pixel in the Guil basin from 2018 September 1st to 2019 June 1st. The pixel elevation is 2026m. Smoothed time series for different η values are shown in different colors. --: NDSI threshold, x: Observations, •: CNI predictions, —: $\text{KS}_{\eta=0.1}$, —: $\text{KS}_{\eta=1}$, —: $\text{KS}_{\eta=10}$, —: $\text{KS}_{\eta=100}$.

by filling only short gaps (up to 3 days) using linear interpolation or by extending the nearest valid measurement. The lines show the NDSI interpolated by the KS for different η values. For $\eta = 0.1$ (yellow), the interpolated time series is close to what would be obtained using simple linear interpolation between observations. This is consistent with the linear nature of the KS, where a small η implies high confidence in the observations. As η increases ($\eta \geq 1$), the signal becomes more heavily regularized. Sharp increases in the NDSI signal, such as those observed in late November or April, are significantly smoothed. In mid-January, fluctuations around the classification threshold cause multiple class changes in the observations, particularly under strong regularization ($\eta = 100$, brown).

Estimation variance grows with longer gaps between observations: the range between minimum and maximum variance narrows, while the average variance increases. This behavior is consistent with high $\eta = \frac{v_{\text{obs}}}{v_{\text{mod}}}$ values, which reflects greater observation noise variance and results in a smaller Kalman gain k_t . Consequently, the observations are given less weight, and the state estimate relies more on the random walk model (Eq. 3).

Since we only have observations as ground truth, we can only assess how much the interpolated time series diverges from the original class distribution. Staying too close to this distribution may indicate that observation noise was fully retained, while diverging too far suggests excessive regularization, removing relevant dynamics from the signal.

5.2.2. Impact of the interpolation methods on models performance

Figure 8 shows the accuracy and F1-scores of models trained with labels interpolated using the KS, for $\eta \in [10^{-3}; 10^2]$. Performance increases with η , reaching its highest value at $\eta = 0.3$.




Figure 7: Snow fraction in the Guil 2018 total set (train, validation and test) for each orbit as a function of the regularization parameter η .

•: Descending 139, •: Ascending 88, •: Descending 66

Set	Interpolation	NoData (%)	Snow (%)	No Snow (%)	f_{Snow}/f_{NoSnow}
Train	None	61.3	19.1	19.6	0.97
	CNI	27.7	39.2	33.1	1.18
	KS $_{\eta=1}$	2.4	51.5	46.1	1.11
	KS $_{\eta=10}$	2.4	49.7	47.9	1.03
Validation	None	27.3	37.7	35.0	1.07
Test	None				
Guil 2018		36.0	37.7	26.3	1.43
Guil 2019		49.28	29.02	21.70	1.34
Gyronde 2018		46.09	33.23	20.68	1.61

Table 5: Class distribution on train, validation and tests sets after projection on SAR geometry for different interpolations.

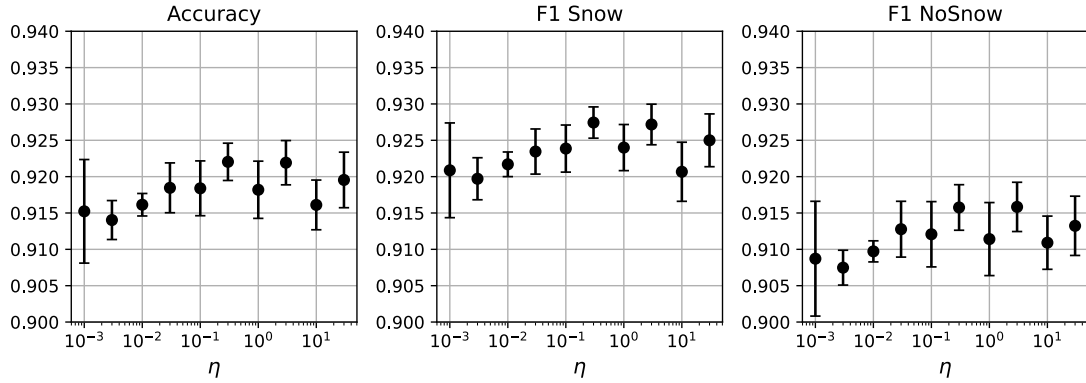


Figure 8: Performance of models trained with KS interpolated labels with different η values on validation set. Mean scores are shown with a dot and standard deviation with a line.

Beyond this point, mean performance trend is unclear and the variance increases. Since mean performance is maximum for $\eta = 0.3$, we select this value as the optimal η for the rest of the

experiments. We refer to models trained without any interpolated labels as $\mathcal{M}_{\text{None}}$, which serve as our baseline. Models trained with CNI-interpolated labels are denoted \mathcal{M}_{CNI} , while \mathcal{M}_{KS} refers to models trained with KS-interpolated labels.

Interpolation	Accuracy \uparrow	Overall F1 \uparrow	F1-Snow \uparrow	F1-No Snow \uparrow
Test Guil 2018 (Same domain)				
None	0.901 (0.0044)	0.902 (0.0042)	0.921 (0.0041)	0.867 (0.0044)
CNI	0.920 (0.0066)	0.920 (0.0064)	0.937 (0.0057)	0.891 (0.0077)
KS	0.927 (0.0046)	0.928 (0.0045)	0.943 (0.0039)	0.901 (0.0055)
Test Guil 2019 (Temporal transfer)				
None	0.904 (0.0037)	0.904 (0.0036)	0.917 (0.0034)	0.886 (0.0042)
CNI	0.923 (0.0093)	0.923 (0.0093)	0.933 (0.0079)	0.908 (0.0113)
KS	0.931 (0.0102)	0.931 (0.0102)	0.939 (0.0088)	0.919 (0.0120)
Test Gyrone 2018 (Spatial transfer)				
None	0.869 (0.0092)	0.871 (0.0089)	0.886 (0.0096)	0.847 (0.0080)
CNI	0.905 (0.0110)	0.906 (0.0107)	0.919 (0.0104)	0.885 (0.0111)
KS	<u>0.905</u> <u>(0.0094)</u>	<u>0.906</u> <u>(0.0091)</u>	<u>0.919</u> <u>(0.0087)</u>	0.884 (0.0098)

Table 6: Mean performance and standard deviations of models trained with pseudo-labels from different interpolation methods on three test sets, using MODIS NDSI labels. The first evaluates without domain shift, the second with temporal domain shift, and the third with spatial domain shift. Best mean scores are shown in bold; when two scores are tied, the one with the lower variance is underlined.

Table 6 shows the evaluation on MODIS NDSI labels on all three test sets. The models are trained on pseudo-labels from each of our interpolation methods: our baseline with no interpolation, CNI and KS with $\eta = 0.3$. On the Guil 2018 test set, label interpolation increases both accuracy and Overall F1-score by approximately 2%, whether CNI or KS is used. \mathcal{M}_{CNI}

has a higher variance, likely due to tendency of CNI to introduce label noise by extending the last valid label in a pixel time series. \mathcal{M}_{KS} yields slightly better average performance than \mathcal{M}_{CNI} , although the gain from no interpolation is smaller in comparison (less than 1% gain).

For temporal transfer, label interpolation also improves model performance. Moreover, all models, including our baseline, achieve higher scores than when tested on the same domain as training, we address this in 6.2.3. KS interpolation on labels still yields better results than CNI. In the Guil 2019 set, every date of the time series is used to evaluate the model instead of a subsample of the 2018 water year, which can explain the increase of performance from the test without transfer.

For spatial transfer, mean performance decrease across all models, with increased associated variance. As for the two previous domains, increasing the amount of pseudo-supervision through label interpolation improves predictions. \mathcal{M}_{CNI} and \mathcal{M}_{KS} yield similar average performance, with \mathcal{M}_{KS} exhibiting slightly lower variance. The interpolation method does not impact model performance, but still improves from the baseline.

Figure 9 shows model predictions on the same domain, with temporal transfer and spatial transfer, respectively. Across all domains, label interpolation leads to smoother predicted snow cover maps and improves recall on the *Snow* class. In the case without transfer shown in Figure 9, \mathcal{M}_{CNI} slightly outperforms \mathcal{M}_{KS} , but the resulting maps are visually similar and closely match the observations. Model predictions on temporal transfer scenarios also lead to snow cover maps close to observations. In the case of spatial transfer however, the predicted snow cover maps are less satisfactory. Models all predict the South-East area of the basin as *NoSnow*, and similar errors persist across multiple dates (not shown in the article). The drop in performance comes from entire patches that are missclassified instead of imprecise segmentation at the interface between classes, unlike in the Guil basin. The main snow-covered areas, however, are still correctly identified.

5.3. Comparison with existing products

5.3.1. THEIA L2B Snow product

The qualitative assessment in Section 5.2.2 shows that our models produce smoother maps that do not exhibit the geometric effect of the projection of MODIS NDSI in the SAR geometry. To evaluate our model prediction at a resolution in the same order of magnitude as the one of the Sentinel-1 image (5×20 m) we also evaluate our models on THEIA L2B Snow product (20×20 m).

We also measure the agreement between thresholded MODIS NDSI and THEIA L2B Snow

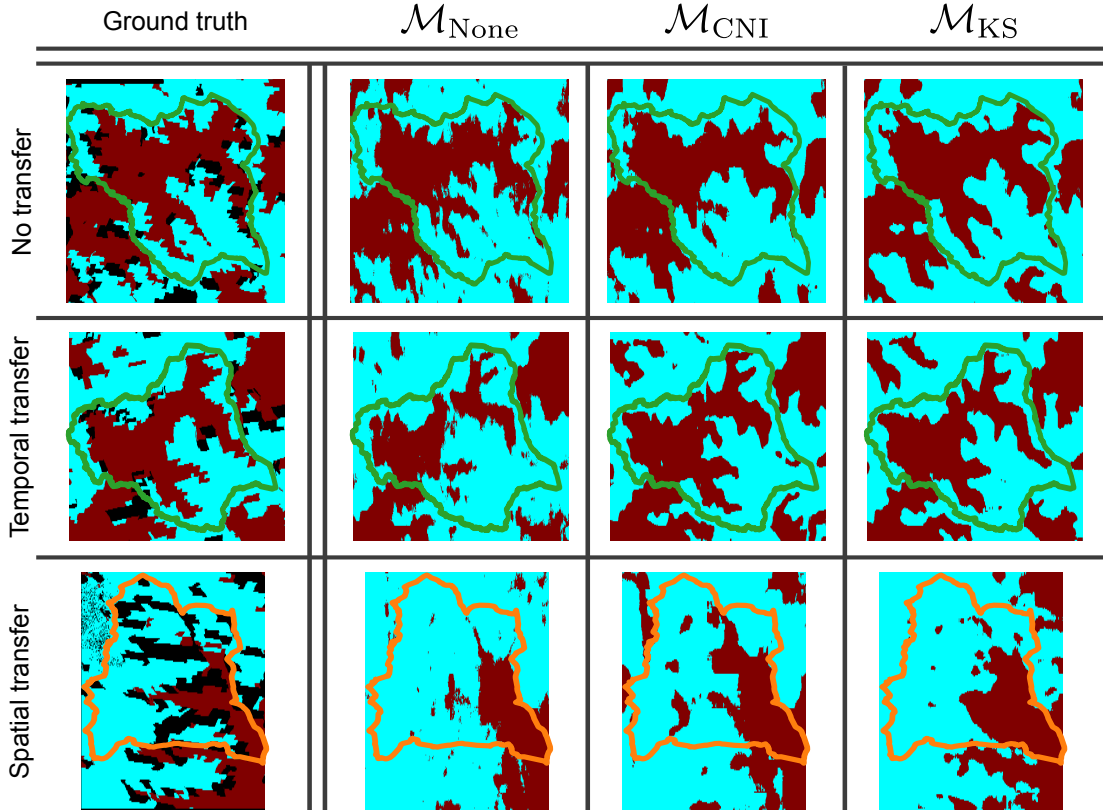


Figure 9: Prediction of models trained with labels from different interpolation method, without domain shift, and under temporal and spatial shift. First row shows Guil basin on 2019 March 31st, second row shows Guil basin on 2020 April 10th and third row Gyronde basin on 2019 March 19th. —: Guil basin, —: Gyronde basin, □: *Snow*, ■: *No Snow*, ■: *NoData*.

to get the baseline performance, summarized in Table 7. MODIS NDSI and THEIA L2B Snow product exhibit a good agreement with an Overall F1-score above 0.88 over the Guil basin for 2018 and 2019 water year. There is less agreement over the Gyronde basin, likely due to a rougher terrain. Figure 10 shows a snow cover map of the same day from MODIS NDSI and THEIA L2B Snow, we see that the small variations in snow cover distribution cannot be captured by MODIS NDSI coarse resolution.

Table 7 also shows model performance for different interpolations on THEIA L2B Snow labels. The interpolation also increases performance, with the model trained on KS pseudo-labels being the best model. However, the performance are lower than the one measured on MODIS NDSI and summarized in Table 6. The Overall F1-score loses 0.065 when evaluated on the Guil 2018. Figure 10 shows that the classification errors are mostly present in the interface between *Snow* and *NoSnow* classes. The model does not manage to predict the granularity at

Interpolation	Accuracy \uparrow	Overall F1 \uparrow	F1-Snow \uparrow	F1-No Snow \uparrow
Test Guil 2018 (Same domain)				
MODIS NDSI	0.896 (0.0000)	0.895 (0.0000)	0.920 (0.0000)	0.850 (0.0000)
None	0.813 (0.0037)	0.815 (0.0033)	0.847 (0.0041)	0.759 (0.0021)
CNI	0.843 (0.0061)	0.843 (0.0053)	0.876 (0.0069)	0.786 (0.0025)
KS	0.864 (0.0031)	0.863 (0.0031)	0.895 (0.0025)	0.808 (0.0053)
Test Guil 2019 (Temporal transfer)				
MODIS NDSI	0.882 (0.0000)	0.881 (0.0000)	0.888 (0.0000)	0.875 (0.0000)
None	0.813 (0.0034)	0.813 (0.0033)	0.825 (0.0040)	0.799 (0.0031)
CNI	0.821 (0.0037)	0.820 (0.0040)	0.842 (0.0042)	0.795 (0.0073)
KS	0.862 (0.0038)	0.861 (0.0038)	0.875 (0.0034)	0.845 (0.0047)
Test Gyronde 2018 (Spatial transfer)				
MODIS NDSI	0.821 (0.0000)	0.816 (0.0000)	0.847 (0.0000)	0.782 (0.0000)
None	0.715 (0.0051)	0.713 (0.0049)	0.744 (0.0085)	0.677 (0.0082)
CNI	0.787 (0.0061)	0.786 (0.0055)	0.810 (0.0088)	0.758 (0.0030)
KS	0.781 (0.0051)	0.779 (0.0049)	0.804 (0.0061)	0.751 (0.0057)

Table 7: Mean performance and standard deviations of models trained with different interpolation methods, evaluated on the THEIA L2B snow product. These labels have a 20m spatial resolution and are derived from Sentinel-2 and Landsat-8 data. For comparison, we also report the scores of MODIS-derived label maps in the first row for each domain. The best model scores are shown in bold.

the interface even if it predicts on a 5×20 m grid, because of the low-resolution training labels.

Even if these metrics are lower than the one computed on MODIS NDSI, the models applied only on SAR dual-pol amplitude images achieve Overall F1-score above 0.86 and are comparable to purely optical based total snow products.

5.3.2. Copernicus Wet/Dry Snow product




Figure 10: March 31st, 2019 snow cover map prediction from \mathcal{M}_{KS} (without domain transfer), with label maps from MODIS NDSI, THEIA L2B Snow and Copernicus WDS, as well as the corresponding label agreement maps. Copernicus WDS does not estimate snow cover in vegetated, water, or urban areas, nor in regions affected by SAR shadow, foreshortening, or layover. Greenish colors indicate correct predictions, while warm colors indicate model errors.

To assess the performance of our method on wet and dry snow independently, we also evaluate it on Copernicus WDS product. Since the models only predict binary snow cover, we compute only the recall for wet and dry snow classes. Tab. 8 shows recalls on both type of snow for models trained with our different interpolation methods. Label interpolation improves recall for both snow types. All models, except \mathcal{M}_{KS} in the spatial transfer setting, achieve higher recall on dry snow than on wet snow. Temporal transfer does not significantly affect performance in either class, while spatial transfer leads to a drop in recall for both wet and dry snow. Figure 10 also includes a label map from Copernicus WDS product, along with a label agreement map. Both types of snow are detected accurately and, as with previous comparisons using MODIS

and THEIA L2B Snow products, most segmentation errors are located at the interface between snow-covered and snow-free areas.

Interpolation	Recall Wet \uparrow		Recall Dry \uparrow	
Test Guil 2018 (Same domain)				
None	0.931	(0.0161)	0.963	(0.0061)
CNI	0.954	(0.0108)	0.971	(0.0041)
KS	0.960	(0.0072)	0.976	(0.0012)
Test Guil 2019 (Temporal transfer)				
None	0.921	(0.0223)	0.974	(0.0043)
CNI	0.955	(0.0090)	0.979	(0.0032)
KS	0.956	(0.0074)	0.975	(0.0051)
Test Gyrone 2018 (Spatial transfer)				
None	0.866	(0.0355)	0.879	(0.0207)
CNI	0.919	(0.0195)	0.927	(0.0124)
KS	0.936	(0.0082)	0.920	(0.0111)

Table 8: Mean recall and standard deviation for wet and dry snow classes, evaluated with Copernicus WDS data.

6. Discussion

6.1. Selection of SAR channels

Quantitative and qualitative evaluation of the different channel sets shows that adding a reference image, as a concatenation or a ratio, always improves model performance. This allows the network to have information, in pixel values, relative to a snow-free state. However, the ratio removes information about elevation, which helps the network to predict less noisy snow cover maps and more accurate interface between snow covered and snow free regions as seen in Figure 5. It seems that the reference image also contains topographical information that is useful for snow cover prediction, and using it for normalization removes this important part. Gallet et al. (2024) showed that the inclusion of topographical information (elevation and slope) for the classification of wet snow patches significantly improved performances. They used ratios between raw backscatter amplitude and reference images, which also improved performance from raw backscatter but did not try a direct concatenation between both. Daudt et al. (2023) also uses elevation maps as input to their RNN. We chose not to use topographical information as we wanted our model to be able to predict snow only from SAR images, and prevent it to overfit on elevation information due to its strong correlation with snow presence. Even if we do not explicitly provide topographical information, the model is able to use the spatial features in the reference image to predict more accurate snow cover maps.

By using snow physics-aware channels to help the network, we have shown that keeping low-frequencies prevents the network to have topographical information and the noisy inputs prevent good snow cover segmentation for the network. As we said in 1, Nagler and Rott (2000) used spatial filtering methods to remove the noise in R_{wet} in order to apply a fixed threshold, however we expected the network to be able to handle this noise and act as an adaptive thresholder in order to detect wet and dry snow. We can note that the reference image is a temporally filtered image, and brings spatial information without taking the risk of removing possibly useful features by filtering the acquisition at time t . R_{wet} is a linear combination of channels set **C**: the additional information in channel set **D** is contained in R_{dry} , which has the purpose to help the network to detect dry snow. However, the only effect seems to make model predictions less stable. In Lievens et al. (2019), the ratio $\sigma_{\text{VH}}^0/\sigma_{\text{VV}}^0$ is computed to see the increase of σ_{VH}^0 due to dry snow accumulation. The ratio with σ_{VV}^0 is used to be independent from temporal variations such as vegetation growth, ground surface changes or snow conditions that impact both σ_{VV}^0 and σ_{VH}^0 . To retrieve snow height, this ratio is compared to the one of the previous day. The difference between those two terms is correlated to snow height variation between the two days. In our application, R_{dry} is computed from one reference image instead of using consecutive observations. The ratio is representative of the depolarization from a snow-free state, but the reference might be too far from observations to be relevant for the model.

6.2. Weak supervision through label interpolation

6.2.1. Setting η value in the Kalman smoother

The KS provides a continuum of interpolated signals through its parametrization with η , ranging from closely following observations (similar to linear interpolation between observations) to producing highly smoothed signals where temporal trends dominate and short-term fluctuations are suppressed. Since these time series are used to assign labels to our pixels, excessive regularization introduces unintended label noise by overly damping oscillations around the classification threshold or by removing potential short-time events. In the opposite case, no regularization keeps the original label noise. In the absence of ground-truth, the optimal level of regularization must be determined empirically. As shown Figure 7, η values smaller than 0.3 produce almost identical time series, leading only to few decision changes in the pseudo-labels. Indeed linear interpolation is a limit case of the Kalman smoothing for small η . However, as shown in Figure 8, these small changes still lead to an increase in model performance for η ranging from $\eta = 0.001$ to $\eta = 0.1$. For larger η values, the performance becomes unstable due to the over smoothing of the NDSI time series. Even if other η values could have lead to higher

performance, we choose to keep $\eta = 0.3$ to ensure the robustness of the model.

6.2.2. Impact of the interpolation methods on the model performance

Increasing the level of supervision through label interpolation, whether using CNI or KS, improved model performance for every task. This improvement was particularly evident in qualitative evaluation, where label interpolation led to smoother predicted snow cover maps. This increased smoothness can be attributed to the gap-filling applied to the training labels: even though the interpolation is purely temporal, spatially close pixels often exhibit strong correlation, resulting in smoother training images. As a result, the network is exposed to more spatially coherent supervision. However, this increased smoothness can also reduce the model ability to capture sharp transitions between snow-covered and snow-free areas, slightly affecting the spatial precision of the predictions. This is especially relevant in SAR imagery, where sharp transitions are expected due to geometric distortions such as layover or foreshortening.

6.2.3. Transfer to other domains

In temporal transfer, all models yield better performance than when tested on the same domain as training, which may seem counter-intuitive. This is explained by the composition of the Guil 2019 test set, which spans the time series from September 1st to June 1st, whereas the Guil 2018 test set covers only a limited subset of the same hydrological period from the previous year (Figure 4). Thus, the 2019 set includes more observations taken during early autumn and mid-winter, when one class tends to dominate. As illustrated in Figure 5, models perform well when predicting large snow covered or snow free areas, while fine-grained segmentation near class boundaries remains more challenging.

The performance gap between predictions in the Guil and Gyrone basins may come from the limited size and diversity of our dataset, which consists of the time series of a single basin, over one water year. While it is enough to train models that perform well on the same domain or under small temporal shifts, the SAR features may not capture enough variability to allow robust generalization to new regions. Our interpolation method does not tackle this issue as it can only increase the amount of optical labels, not diversify the SAR features.

6.3. Comparison with existing products

We validated our method using the THEIA L2B Snow product, which is derived from Sentinel-2 and Landsat-8 optical imagery and maps total snow cover. While it provides higher spatial resolution than MODIS, it is only available every 5 to 10 days, but have been all included in the Guil 2018 test set.

The evaluation on THEIA L2B Snow product, listed in Table 7, confirms the gain of using pseudo-label obtain by temporal interpolation of the snow product during training. However, the comparison with the Overall F1-score evaluated on the MODIS NDSI product listed in 6, shows that all the models is lower when evaluated on the THEIA L2B Snow product than evaluated on the MODIS NDIS.

To evaluate how this loss can be linked to the use of coarser resolution label during training, we performed the direct comparison between the thresholded MODIS NDSI and the THEIA L2B Snow product. Although the agreement between the two products is high, the finer resolution of the THEIA L2B Snow product reveals small-scale features that are not captured in the coarser resolution MODIS NDSI. Thus, most disagreements are located at the boundaries between the *Snow* and *No Snow* classes, particularly at mid-altitudes, where rapid spatial changes in snow cover require higher-resolution observations. This is especially true for the Gyrone basin where the disagreement between the MODIS NDSI and the THEIA L2B Snow product is the highest.

The model trained using the Kalman smoothed pseudo-label achieves performance close to the agreement between the two products, showing its ability to predict total snow cover from SAR dual-polarization amplitude images. Unlike MODIS NDSI, it produces spatially smooth total snow cover maps, showing its capacity to build on small scale SAR features. However, classification errors are still concentrated at class boundaries, as shown in Figure 10. In these transition areas, small-scale patterns are often missed, likely due to the coarse pseudo-labels derived from MODIS NDSI used during training.

To ensure the deep-learning models do not make errors on a single type of snow, we compared the estimated snow cover maps to the Copernicus WDS product, which distinguishes between wet and dry snow. The comparison showed that both snow classes were correctly detected, but with more errors on the *wet Snow* class. This is somewhat counterintuitive, since dry snow is largely transparent to SAR signals.

However, Blöschl et al. (1990) showed that in a basin, the snowpack at the lowest altitudes is saturated with liquid water and runoff-producing. The model lower recall on the *Wet Snow* class could be explained by its spatial repartition and the model difficulties to segment accurately *Snow* and *No Snow* classes boundaries, as shown in Figure 10.

The small proportion of *Wet Snow* in the training set could also explain the model difficulties to segment this class. In the MODIS NDSI product from the 2019 Guil test, we observed a count of *Dry Snow* pixels twice higher than the count of *Wet Snow* pixels among the pixels classified as *Snow*. Because Sentinel-1 and Sentinel-2 have regular revisit time, we can assume this proportion

is representative of the one in the Guil 2018 train set. An under-representation of the *Wet Snow* class in the positive examples given to the model during training could also explain why its associated recall is lower.

7. Conclusion

In this study, we proposed a method to predict binary total snow cover maps from Sentinel-1 SAR backscatter amplitude. Training labels were generated in an unsupervised manner by thresholding MODIS NDSI time series from the MOD10A1 dataset. A U-Net architecture was trained for binary semantic segmentation using a standard cross-entropy loss.

We showed that the most effective SAR input configuration is the concatenation of co- and cross-polarized backscatter amplitude channels with reference images derived from snow-free acquisitions. More heavily processed, snow-specific SAR channels did not improve performance, likely due to the loss of topographic information that the network otherwise learns to exploit.

We investigated whether adding pseudo-labels derived from MODIS NDSI time series interpolation helps the model to predict better snow cover maps. We applied temporal interpolation on MODIS NDSI time series using two approaches: Cloest Neighbours Interpolation and a Kalman smoother. The level of regularization in the KS was found to be critical: moderate values improved performance by balancing temporal consistency and fidelity to observations, whereas excessive smoothing led to information loss. We have shown that models trained with interpolated labels achieved better accuracy and produced less noisy predictions. Under temporal domain shift, the models were able to generalize better when trained with pseudo-labels. However, spatial generalization remains a challenge, as the training set may lack diversity in SAR features to capture the variability across different basins. Increasing label density cannot fully address this limitation.

These results highlight that total snow cover maps can be estimated solely from SAR dual-pol amplitude images by leveraging the ability of optical data to detect both dry and wet snow to supervise model training, and that adding pseudo-labels with temporal interpolation improves prediction in cloud-prone environment.

The commonly used 0.4 threshold may not be optimal under certain observation conditions, such as shadow or forest cover, which can lead to underestimation of snow-covered pixels (Härer et al., 2018; Zhang et al., 2019). Adding label uncertainty knowledge during training could improve performance and remove the need of using a fixed threshold. Finally, we aim to extend these methods to other land cover types, such as vegetation for deforestation monitoring, which

show distinct responses in both SAR and optical signals, akin to snow.

References

- Blöschl, G., Kirnbauer, R., Gutknecht, D., 1990. Modelling snowmelt in a mountainous river basin on an event basis. *Journal of Hydrology* 113, 207–229. URL: [http://dx.doi.org/10.1016/0022-1694\(90\)90176-X](http://dx.doi.org/10.1016/0022-1694(90)90176-X), doi:10.1016/0022-1694(90)90176-x.
- Brun, E., Martin, E., Simon, V., Gendre, C., Coleou, C., 1989. An energy and mass model of snow cover suitable for operational avalanche forecasting. *Journal of Glaciology* 35, 333–342. URL: <http://dx.doi.org/10.3189/S0022143000009254>, doi:10.3189/S0022143000009254.
- Daudt, R.C., Wulf, H., Hafner, E.D., Bühler, Y., Schindler, K., Wegner, J.D., 2023. Snow depth estimation at country-scale with high spatial and temporal resolution. *ISPRS Journal of Photogrammetry and Remote Sensing* 197, 105–121. doi:10.1016/j.isprsjprs.2023.01.017.
- Deng, G., Tang, Z., Dong, C., Shao, D., Wang, X., 2024. Development and evaluation of a cloud-gap-filled MODIS Normalized Difference Snow Index product over high mountain Asia. *Remote Sensing* 16, 192. URL: <http://dx.doi.org/10.3390/rs16010192>, doi:10.3390/rs16010192.
- Dozier, J., 1984. Snow reflectance from LANDSAT-4 thematic mapper. *IEEE Trans. Geosci. Remote Sens.* GE-22, 323–328.
- Dozier, J., 1989. Spectral signature of alpine snow cover from the landsat thematic mapper. *Remote Sensing of Environment* 28, 9–22. URL: [http://dx.doi.org/10.1016/0034-4257\(89\)90101-6](http://dx.doi.org/10.1016/0034-4257(89)90101-6), doi:10.1016/0034-4257(89)90101-6.
- Erxleben, J., Elder, K., Davis, R., 2002. Comparison of spatial interpolation methods for estimating snow distribution in the Colorado Rocky Mountains. *Hydrological Processes* 16, 3627–3649. URL: <http://dx.doi.org/10.1002/hyp.1239>, doi:10.1002/hyp.1239.
- Gallet, M., Atto, A., Karbou, F., Trouvé, E., 2024. Wet snow detection from satellite sar images by machine learning with physical snowpack model labeling. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 17, 2901–2917. doi:10.1109/JSTARS.2023.3342990.

- Gascoïn, S., Barrou Dumont, Z., Deschamps-Berger, C., Marti, F., Salgues, G., López-Moreno, J.I., Revuelto, J., Michon, T., Schattan, P., Hagolle, O., 2020. Estimating Fractional Snow Cover in open terrain from Sentinel-2 using the Normalized Difference Snow Index. *Remote Sensing* 12, 2904. URL: <http://dx.doi.org/10.3390/rs12182904>, doi:10.3390/rs12182904.
- Global Climate Observing System, 2006. Systematic observation requirements for satellite-based products for climate: supplemental details to the satellite-based component of the Implementation Plan for the Global Observing System for Climate in support of the UNFCCC. Technical Report GCOS-107 (WMO/TD No. 1338). World Meteorological Organization. URL: <https://www.globsnow.info/docs/gcos-107.pdf>.
- Hall, D., Riggs, G., 2021. [dataset]MODIS/Terra Snow Cover Daily L3 Global 500m SIN Grid, Version 61. URL: <http://nsidc.org/data/MOD10A1/versions/61>, doi:10.5067/MODIS/MOD10A1.061.
- Hall, D.K., Riggs, G.A., Salomonson, V.V., 1995. Development of methods for mapping global snow cover using Moderate-Resolution Imaging Spectroradiometer data. *Remote Sensing of Environment* 54, 127–140. URL: <https://www.sciencedirect.com/science/article/pii/S003442579500137P>, doi:[https://doi.org/10.1016/0034-4257\(95\)00137-P](https://doi.org/10.1016/0034-4257(95)00137-P).
- Hall, D.K., Riggs, G.A., Salomonson, V.V., DiGirolamo, N.E., Roman, M.O., 2021. [dataset]MODIS/Terra Snow Cover Daily L3 Global 500m SIN Grid, Version 6 (MOD10A1F). URL: <https://doi.org/10.5067/MODIS/MOD10A1F.061>, doi:10.5067/MODIS/MOD10A1F.061.
- Härer, S., Bernhardt, M., Siebers, M., Schulz, K., 2018. On the need for a time- and location-dependent estimation of the NDSI threshold value for reducing existing uncertainties in snow cover maps at different scales. *Cryosphere* 12, 1629–1642.
- He, G., Feng, X., Xiao, P., Xia, Z., Wang, Z., Chen, H., Li, H., Guo, J., 2017. Dry and wet snow cover mapping in mountain areas using SAR and optical remote sensing data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 10, 2575–2588. URL: <http://dx.doi.org/10.1109/JSTARS.2017.2673409>, doi:10.1109/jstars.2017.2673409.
- Hosang, J., Dettwiler, K., 1991. Evaluation of a water equivalent of snow cover map in a small catchment area using a geostatistical approach. *Hydrological Processes* 5, 283–290. URL: <http://dx.doi.org/10.1002/hyp.3360050308>, doi:10.1002/hyp.3360050308.

- Hou, J., Huang, C., Zhang, Y., Guo, J., Gu, J., 2019. Gap-filling of MODIS fractional snow cover products via non-local spatio-temporal filtering based on machine learning techniques. URL: <http://dx.doi.org/10.3390/rs11010090>, doi:10.3390/rs11010090.
- Immerzeel, W.W., Lutz, A.F., Andrade, M., Bahl, A., Biemans, H., Bolch, T., Hyde, S., Brumby, S., Davies, B.J., Elmore, A.C., Emmer, A., Feng, M., Fernández, A., Haritashya, U., Kargel, J.S., Koppes, M., Kraaijenbrink, P.D.A., Kulkarni, A.V., Mayewski, P.A., Nepal, S., Pacheco, P., Painter, T.H., Pellicciotti, F., Rajaram, H., Rupper, S., Sinisalo, A., Shrestha, A.B., Viviroli, D., Wada, Y., Xiao, C., Yao, T., Baillie, J.E.M., 2019. Importance and vulnerability of the world's water towers. *Nature* 577, 364–369. URL: <http://dx.doi.org/10.1038/s41586-019-1822-y>, doi:10.1038/s41586-019-1822-y.
- Institut national de l'information géographique et forestière (IGN), 2021. [dataset]RGE ALTI® 1 m - Modèle numérique de terrain. <https://geoservices.ign.fr/documentation/donnees/rgealti.html>. Licence Etalab v2.0.
- James, G., Karbou, F., Durand, P., 2024. A dynamical system approach to wet snow retrieval using Sentinel-1 SAR images, in: IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium, IEEE. p. 115–120. URL: <http://dx.doi.org/10.1109/IGARSS53475.2024.10641229>, doi:10.1109/igarss53475.2024.10641229.
- Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering* 82, 35–45.
- Karbou, F., Veyssi re, G., Coleou, C., Dufour, A., Gouttevin, I., Durand, P., Gascoin, S., Grizonnet, M., 2021. Monitoring wet snow cover an alpine region using Sentinel-1 observations. *Remote Sensing* 13, 381. doi:10.3390/rs13030381.
- L e, T.T., Atto, A., Trouv e, E., Karbou, F., 2023. Deep semantic fusion of Sentinel-1 and Sentinel-2 snow products for snow monitoring in mountainous regions, in: IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium, IEEE, Pasadena, CA, USA. pp. 6286–6289. doi:10.1109/IGARSS52108.2023.10282065.
- Lemke, P., Ren, J., Alley, R.B., Allison, I., Carrasco, J., Flato, G., Fujii, Y., Kaser, G., Mote, P., Thomas, R.H., et al., 2007. Observations. Changes in snow, ice and frozen ground. Chapter 4 .

- Li, M., Zhu, X., Li, N., Pan, Y., 2020. Gap-filling of a MODIS Normalized Difference Snow Index product based on the similar pixel selecting algorithm: a case study on the Qinghai–Tibetan plateau. *Remote Sensing* 12, 1077. URL: <http://dx.doi.org/10.3390/rs12071077>, doi:10.3390/rs12071077.
- Lievens, H., Demuzere, M., Marshall, H.P., Reichle, R.H., Brucker, L., Brangers, I., de Rosnay, P., Dumont, M., Girotto, M., Immerzeel, W.W., Jonas, T., Kim, E.J., Koch, I., Marty, C., Saloranta, T., Schöber, J., De Lannoy, G.J.M., 2019. Snow depth variability in the northern hemisphere mountains observed from space. *Nat. Commun.* 10, 4629.
- Man, Z., Xie, C., Jiang, R., Che, S., 2022. Freeze–thaw cycle frequency affects root growth of alpine meadow through changing soil moisture and nutrients. *Scientific Reports* 12. URL: <http://dx.doi.org/10.1038/s41598-022-08500-w>, doi:10.1038/s41598-022-08500-w.
- Mastrotheodoros, T., Pappas, C., Molnar, P., Burlando, P., Manoli, G., Parajka, J., Rigon, R., Szeles, B., Bottazzi, M., Hadjidoukas, P., Fatichi, S., 2020. More green and less blue water in the Alps during warmer summers. *Nature Climate Change* 10, 155–161. URL: <http://dx.doi.org/10.1038/s41558-019-0676-5>, doi:10.1038/s41558-019-0676-5.
- Montginoux, M., Weissgerber, F., Lobry, S., Idier, J., 2023. Évaluation du couvert neigeux à partir d’images SAR par apprentissage profond basé sur des images optiques de référence, in: 29e colloque GRETSI, Grenoble (38000), France. URL: <https://hal.science/hal-04256105>.
- Moreno-Martínez, Á., Izquierdo-Verdiguier, E., Maneta, M.P., Camps-Valls, G., Robinson, N., Muñoz-Marí, J., Sedano, F., Clinton, N., Running, S.W., 2020. Multispectral high resolution sensor fusion for smoothing and gap-filling in the cloud. *Remote Sensing of Environment* 247, 111901. URL: <http://dx.doi.org/10.1016/j.rse.2020.111901>, doi:10.1016/j.rse.2020.111901.
- Nagler, T., Rott, H., 2000. Retrieval of wet snow by means of multitemporal sar data. *IEEE Transactions on Geoscience and Remote Sensing* 38, 754–765. doi:10.1109/36.842004.
- Nagler, T., Rott, H., Ripper, E., Bippus, G., Hetzenecker, M., 2016. Advancements for snowmelt monitoring by means of Sentinel-1 SAR. *Remote Sensing* 8, 348. URL: <https://www.mdpi.com/2072-4292/8/4/348>, doi:10.3390/rs8040348.
- Phillips, D.L., Dolph, J., Marks, D., 1992. A comparison of geostatistical procedures for spatial analysis of precipitation in mountainous terrain. *Agricultural and Forest Meteorology*

- 58, 119–141. URL: [http://dx.doi.org/10.1016/0168-1923\(92\)90114-J](http://dx.doi.org/10.1016/0168-1923(92)90114-J), doi:10.1016/0168-1923(92)90114-j.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer.
- Rouhier, L., 2018. Régionalisation d'un modèle hydrologique distribué pour la modélisation de bassins non jaugés. Application aux vallées de la Loire et de la Durance. Ph.D. thesis. Sorbonne Université.
- Salomonson, V., Appel, I., 2004. Estimating fractional snow cover from MODIS using the Normalized Difference Snow Index. *Remote Sensing of Environment* 89, 351–360. URL: <https://www.sciencedirect.com/science/article/pii/S0034425703002864>, doi:<https://doi.org/10.1016/j.rse.2003.10.016>.
- Sedano, F., Kempeneers, P., Hurtt, G., 2014. A Kalman filter-based method to generate continuous time series of medium-resolution NDVI images. *Remote Sensing* 6, 12381–12408. URL: <https://www.mdpi.com/2072-4292/6/12/12381>, doi:10.3390/rs61212381.
- Simon, G., Manuel, G., Tristan, K., Germain, S., 2018. Algorithm theoretical basis documentation for an operational snow cover product from Sentinel-2 and Landsat-8 data (Let-it-snow). URL: <https://doi.org/10.5281/zenodo.1414452>, doi:10.5281/zenodo.1414452.
- Singh, G., Venkataraman, G., Rao, Y.S., Kumar, V., Snehamani, 2008. InSAR coherence measurement techniques for snow cover mapping in Himalayan region, in: IGARSS 2008 - 2008 IEEE International Geoscience and Remote Sensing Symposium, IEEE. pp. IV–1077–IV–1080. URL: <http://dx.doi.org/10.1109/IGARSS.2008.4779913>, doi:10.1109/igarss.2008.4779913.
- Strozzi, T., Wegmuller, U., Matzler, C., 1999. Mapping wet snowcovers with SAR interferometry. *International Journal of Remote Sensing* 20, 2395–2403. URL: <http://dx.doi.org/10.1080/014311699212083>, doi:10.1080/014311699212083.
- Tarquini, S., Isola, I., Favalli, M., Battistini, A., 2007. [dataset]TINITALY, a digital elevation model of Italy with a 10 meters cell size, version 1.0. URL: <https://tinitaly.pi.ingv.it/>, doi:10.13127/TINITALY/1.0.

- Teleubay, Z., Yermekov, F., Toleubekova, Z., Shmatov, B., Raiev, Y., Assylkhanova, A., 2023. Snow height and Snow Water Equivalent estimation from Snow Cover Fraction using Sentinel-2 satellite images in North Kazakhstan. *Athens Journal of Sciences* 10, 53–64. URL: <http://dx.doi.org/10.30958/ajs.10-1-4>, doi:10.30958/ajs.10-1-4.
- Tsai, Y.L., Dietz, A., Oppelt, N., Kuenzer, C., 2019. Wet and dry snow detection using Sentinel-1 SAR data for mountainous areas with a machine learning technique. *Remote Sensing* 11, 895. URL: <http://dx.doi.org/10.3390/rs11080895>, doi:10.3390/rs11080895.
- Tsardanidis, I., Koukos, A., Sitokonstantinou, V., Drivas, T., Kontoes, C., 2024. Cloud gap-filling with deep learning for improved grassland monitoring. URL: <https://arxiv.org/abs/2403.09554>, doi:10.48550/ARXIV.2403.09554.
- Viviroli, D., Dürr, H.H., Messerli, B., Meybeck, M., Weingartner, R., 2007. Mountains of the world, water towers for humanity: typology, mapping, and global significance. *Water Resources Research* 43. URL: <http://dx.doi.org/10.1029/2006WR005653>, doi:10.1029/2006wr005653.
- Wang, Y., Wu, W., Zhang, Z., Li, Z., Zhang, F., Xin, Q., 2025. A temporal attention-based multi-scale generative adversarial network to fill gaps in time series of MODIS data for land surface phenology extraction. *Remote Sensing of Environment* 318, 114546. URL: <http://dx.doi.org/10.1016/j.rse.2024.114546>, doi:10.1016/j.rse.2024.114546.
- Weissgerber, F., Charrier, L., Thomas, C., Nicolas, J.M., Trouvé, E., 2022. LabSAR, a one-GCP coregistration tool for SAR–InSAR local analysis in high-mountain regions. *Frontiers in Remote Sensing* 3, 935137. URL: <https://www.frontiersin.org/articles/10.3389/frsen.2022.935137/full>, doi:10.3389/frsen.2022.935137.
- Zhang, H., Zhang, F., Zhang, G., Che, T., Yan, W., Ye, M., Ma, N., 2019. Ground-based evaluation of MODIS snow cover product V6 across china: Implications for the selection of NDSI threshold. *Science of The Total Environment* 651, 2712–2726. URL: <http://dx.doi.org/10.1016/j.scitotenv.2018.10.128>, doi:10.1016/j.scitotenv.2018.10.128.