Automating glacier facies classification: pan-European dataset and deep learning baseline

Konstantin A. Maslov[a,*], Thomas Schellenberger[b], Prashant Pandit[c,d], Claudio Persello[a], Alfred Stein[a]

[a] *University of Twente, Drienerlolaan 5, Enschede, 7522NB, Overijssel, The Netherlands*
[b] *University of Oslo, Sem Sælands vei 1, Oslo, 0371, Østlandet, Norway*
[c] *University of Trento, via Calepina 14, Trento, 38122, Trentino, Italy*
[d] *Eurac Research, Drususallee 1/Viale Druso 1, Bolzano, 39100, South Tyrol, Italy*

Abstract

Glacier facies play a critical role in understanding the mass balance of glaciers, offering insights into accumulation and melting processes. Accurate mapping of glacier facies is therefore essential for monitoring glacier response to climate change and informing climate policies. In this study, we present the largest glacier facies dataset ever compiled for Europe, comprising 31 glaciers, 92 Landsat and Sentinel-2 scenes, 138 273 expert point labels and eight classes—five glacier facies (*ice*, *snow*, *debris*, *firn* and *superimposed ice*) and three miscellaneous classes (*shadow*, *cloud* and *water*)—encompassing a wide variety of surface conditions. A confident learning method pruned 16% of ambiguous expert labels overall. A compact and straightforward convolutional neural network reached a macro-average F_1 score of 78% on the complete cleaned data or 69% on the full, unpruned data, and 79.9±11.5% glacier-wise. When the facies products were regressed against World Glacier Monitoring Service records, they showed moderate, yet significant correlation with the surface mass balance measurements globally (r = 0.63, RMSE = 0.61 m w.e., where m w.e. denotes metre of water equivalent) and competitive correspondence for glacier-specific calibration (r = 0.81, RMSE = 0.26 m w.e.). Overall, the dataset and baseline show that large-scale glacier facies classification can be achieved with high consistency. By providing both the dataset and baseline classification models, we aim to support the broader community in developing more advanced methods for glacier facies mapping to enhance our understanding of ongoing glacial changes.

Keywords: Glacier mapping, glacier facies, confident learning, deep learning, convolutional neural network, glacier surface mass balance

1. Introduction

Glaciers are among the most sensitive indicators of climate change, with alterations in their extent and mass balance cited as clear evidence of global warming (Rounce et al., 2023; Zemp et al., 2025). Across Europe, glacierised regions exhibit diverse characteristics linked to local climatology, topography and geological features. Monitoring

*k.a.maslov@utwente.nl

these glaciers is essential for understanding regional water resources (Fox et al., 2024), sea-level rise contributions (Zemp et al., 2019) and potential hazards such as glacial lake outburst floods (Carrivick and Tweed, 2016). As glaciers respond to seasonal and long-term changes in temperature and precipitation, the spatial distribution of glacier facies—zones of different snow and ice properties—offer valuable insights into the accumulation and ablation processes.

Maps of glacier snow serve as critical inputs to mass balance studies. Glacier snow cover evolution over time and distribution can be used to infer surface mass balance (SMB) (Rabatel et al., 2017) by means of empirical methods, such as the snow line altitude (SLA) method (Rabatel et al., 2005), the accumulation area ratio (AAR) method (Hock et al., 2007) as well as approaches that rely on relating snow-covered area elevation to SMB (Drolon et al., 2016). Similarly, firn changes correlate with long-term trends in SMB (König et al., 2004). Superimposed ice slows mass balance response to warming, and its maps can serve as a proxy to quantify refreezing within the snowpack (Wright et al., 2005). Supraglacial debris and lakes govern radiation transfer and ablation rates, while the lakes can also influence glacier dynamics by supplying water to the glacier bed (Huo et al., 2021; Wendleder et al., 2021). Although not a facies per se, tracking cast shadows over time allows inferring changes of glacier surface elevation (Pfau et al., 2023). Beyond that, glacier facies maps can potentially be used to refine the parameterisations in physical models, thus, improving the accuracy and reliability of simulations that aim to project future glacier evolution. Hence, mapping all these classes is particularly valuable for improving glacier SMB models, as each provides complementary constraints on accumulation, ablation, refreezing and long-term storage.

The concept of glacier facies was introduced by Müller (1962) through in situ studies on Canadian glaciers, identifying spatially coherent zones of accumulation, refreezing and ablation. With the advent of satellite remote sensing, multispectral imagery enabled facies identification by visual interpretation and simple band/ratio thresholding (Hall et al., 1987, 1988; Williams et al., 1991; Gupta et al., 2005), refined by comparisons with in situ spectra that highlighted the importance of visible and near-infrared (NIR) wavelengths for separating snow and ice facies and shortwave-infrared (SWIR) for detecting debris (Pope and Rees, 2014). These ideas evolved into automated classifiers that exploit both optical and thermal signals as well as topographic context (Bhardwaj et al., 2015; Shukla and Yousuf, 2017). At regional to global scales, related approaches have been used to track facies change, map supraglacial debris and retrieve snowline altitudes for large glacier populations, often combined with manual corrections (Ali et al., 2017; Herreid and Pellicciotti, 2020; Larocca et al., 2024; Loibl et al., 2025). Although wet snow, firn and bare ice can appear spectrally similar in raw imagery, multi-band analysis and inclusion of spatial context have shown that these facies are separable (Dietz et al., 2011; Yousuf et al., 2019; Florath et al., 2021), especially when elevation data are used. Still, delineating subtle transitions remains challenging, underscoring the need for more sophisticated methods.

Recent advances in machine learning, particularly convolutional neural networks (CNNs), have made the automated classification of remotely sensed images increasingly feasible. Machine learning classifiers have been widely adopted for glacier facies classification at sub-regional scales with medium-resolution (Zhang et al., 2019; Florath et al., 2021; Lu et al., 2021; Prieur et al., 2022) and high-resolution (Luis and Singh, 2020; Jawak et al., 2022) optical imagery. Despite these advances, large-scale and generalisable solutions are missing. Most prior efforts either focus

on single-class identification (e.g. snow or debris) or rely on limited, unpublished datasets and do not rigorously test transferability to unseen glaciers or regions. In practice, the heterogeneity of glacier surfaces means that methods tuned to one context fail in another. Notably, Aberle et al. (2025) released a multi-glacier facies dataset focusing on six glaciers in Alaska, with two glaciers used exclusively for validation. Robust mapping requires assembling extensive data that capture a wide diversity of glaciers. At present, the community lacks open benchmarks for developing and evaluating models at scale.

In this study, we address the gap in large-scale, high-quality training data by presenting the largest curated glacier facies dataset assembled for European glaciers. Covering 31 glaciers, our dataset reflects a wide range of climatic and morphological conditions. We systematically curated annotations using a confident learning method (Northcutt et al., 2021) to ensure label reliability. Using this dataset, we trained a simple CNN to provide a benchmark classification performance, thus setting the stage for future methodological improvements. This work contributes two main novelties: (1) a comprehensive dataset for European glacier facies, and (2) the first benchmark results using a lightweight CNN model. By releasing both our dataset and baseline code, we aim to further facilitate advancements in automated glacier facies mapping and enhance our collective understanding of glacier evolution under a changing climate.

2. Study area and data

Our dataset covers 31 glaciers distributed across five European regions—the Pyrenees, the Alps, Scandinavia, Iceland and Svalbard—thereby representing a pan-European selection of glaciers (Figure 1). Specifically, we selected the following glaciers: Gebroulaz, Sex Rouge, Silvretta, Midtre Lovenbreen, Engabreen, Maladeta, Lupo, Corbassiere, Goldberg K, Tsanfleuron, Bruarjokull, Schwarzberg, Graasubreen, Langjokull ice cap, Austre Lovenbreen, Fridtjovbreen, Storbreen, Kongsvegen, Austre Okstindbreen, Kesselwand F, Marmaglaciaeren, Argentiere, Svelgjabreen, Waldemarbreen, Gietro, Grosser Aletsch, Adler, Jamtal F, Skeidararjoekull, Pasterze and Kronebreen. This selection ensures a wide coverage of glacier sizes, elevations and latitudes, capturing the complexity of different climatic conditions, and aligns well with the available SMB records from the World Glacier Monitoring Service (WGMS) (WGMS, 2021).

As input to the classification models, we used a combination of optical reflectance bands together with topographic and geometric features, including elevation, slope, hillshade, shadow mask and proximity to the glacier boundary, to capture both spectral and spatial characteristics of glacier surfaces. In total, we compiled 92 optical images acquired by the Landsat missions (5, 7, 8 and 9) and Sentinel-2 with acquisition dates ranging from 1986 to 2024, with the most acquisitions coming from the end of the ablation season. Some glaciers appear multiple times, enabling a temporal perspective that captures different interannual conditions. We used top-of-atmosphere optical data from blue, green, red, near-infrared and two shortwave-infrared channels consistent among the sensors. Beyond these optical observations, we integrated additional information from the Copernicus Digital Elevation Model (DEM). We derived elevation (normalised per glacier to mitigate glacier-specific elevation ranges) and slope, along with a shadow mask

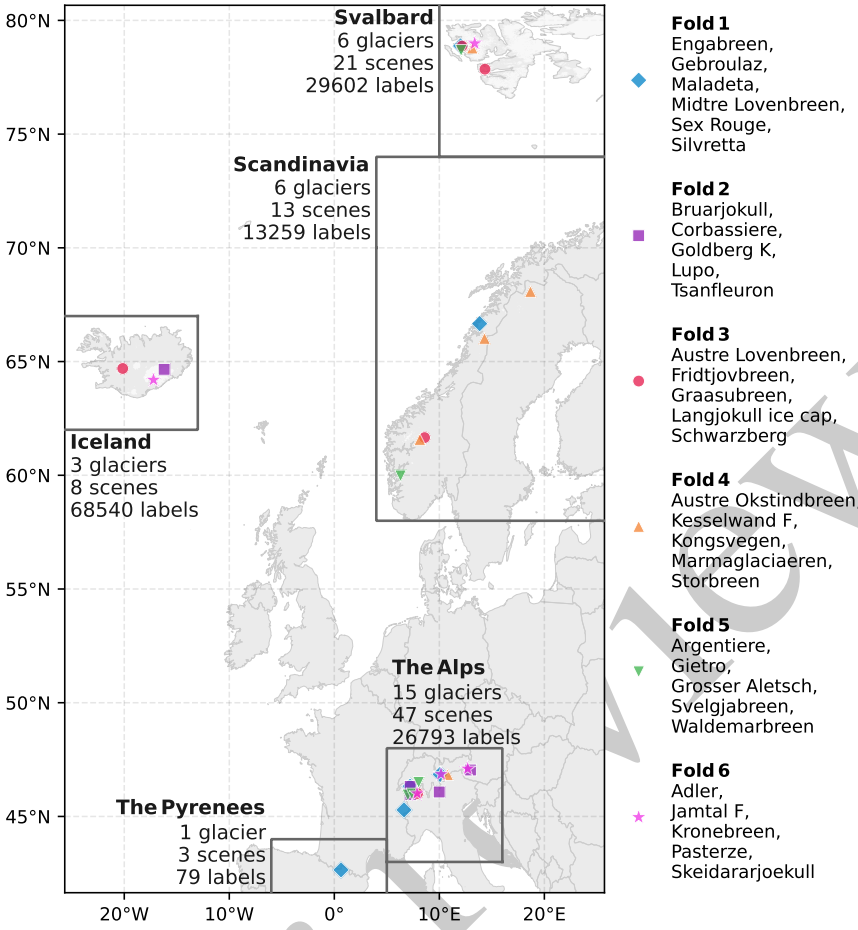

Figure 1: Geographical span and overview of the compiled dataset. Colours denote six glacier folds used for cross-validation (see Section 3).

and a hillshade map. The shadow mask was generated using the `python-dem-shadows` package (de Ruijter, 2016), which allows shadow casting based on the sun azimuth and zenith angles extracted from the optical scene metadata. Similarly, the hillshade (a relief image that shows how brightly each DEM pixel is lit ignoring cast shadows) was created with the same solar geometry. Finally, to provide insight into the relative distance from the glacier perimeter, we computed proximity, i.e. distance to the glacier boundary, and also normalised it per glacier. The glacier outlines were taken directly from the Global Land Ice Measurements from Space database (GLIMS Consortium, 2015) and manually corrected where necessary, e.g., in cases of missing debris-covered ice patches, or generated with GlaViTU (Maslov et al., 2025) for dates without database entries. The GlaViTU-derived outlines were checked visually and found to be of high quality. For subsequent experiments, we organised the feature sets as follows:

1. Optical only
2. Optical + elevation + slope
3. Optical + elevation + slope + shadow mask + hillshade
4. Optical + elevation + slope + shadow mask + hillshade + proximity

Table 1: Pre-cross-check expert agreement and post-cross-check distribution shifts.

	ice	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Total
Expert agreement, %	99.96	80.80	99.04	91.76	99.96	96.22	92.70	84.26	93.61
Distribution shift, %	+3.92	-19.05	-0.61	+36.16	+0.11	-0.13	-2.91	-14.42	+0.19

All images were resampled to 10 m × 10 m grids with bilinear interpolation to match the spatial resolution of the Sentinel-2 data.

Three expert glaciologists manually labelled eight classes observable within the glacier outlines in the satellite imagery. Five are glacier facies: *ice*, *snow*, *debris*, *firn* and *superimposed ice*. *Ice* refers to exposed bare glacier ice free of snow or debris. *Snow* is seasonal accumulation from the current hydrological year. *Firn* is older snow that has survived at least one ablation season and begun densification. *Superimposed ice* denotes refrozen meltwater forming a transparent ice layer on ice or firn. *Debris* represents supraglacial rock or sediment cover. The remaining three classes are not glacier facies: *shadow* (areas affected by topographic shading), *cloud* (any identifiable cloud cover) and *water* (supraglacial and proglacial lakes).

Labelling was conducted on each optical scene, focusing on the glacier surfaces. Typical spectral thresholds served as preliminary guidance for the main classes: pixels with NIR reflectance $\gtrapprox$ 0.55 were labelled *snow*, $\approx$ 0.40–0.55 as *firn*, and $\lessapprox$ 0.40 as *ice*. These limits were applied flexibly, i.e. continuous, visually homogeneous zones (e.g. smooth high-albedo patches) were kept consistent within identifiable snow/ice layers even when individual pixels deviated due to topographical shading. *Superimposed ice* was recognised by its distinctive bright-blue appearance in visible part of the spectrum (R-G-B) and deep-blue shades in SWIR-NIR-red composites. Additional cues included yellow-brown appearance for *firn*, pronounced crevasse texture for *ice*, reddish SWIR-NIR-red tones for *debris* (hence, predominantly attributing mixed ice/debris pixels to *ice*), enhanced SWIR brightness for semi-transparent *cloud*, and clear radiometric minima for *shadow* and *water*. Elevation and contextual signals supported the labelling process. *Snow* mainly occupied higher altitudes. *Firn* was clustered between *ice* and *snow* at the end of the ablation season. *Debris* was typically located at the termini or followed the ice flow as medial moraines if it originated from nunataks. Patches of *superimposed ice* were adjacent to *firn*. *Shadow* outlines followed the topography patterns. The presence of meltwater channels helped to identify *water* bodies. Multi-date imagery, very-high-resolution images and field photographs were consulted whenever available, e.g., from public remote sensing platforms and documented field observations. In cases of high ambiguity, the areas were left unlabelled.

One cross-check round was conducted to correct potential discrepancies between the experts. Most expert disagreement occurred in transitions between *snow*, *firn* and *ice* (Figure 2a,b), where we unified the interpretation as much as possible. After the cross-check, the labelling of *debris* was more consistent across scenes, attributing mixed

debris-ice pixels to *ice* (Figure 2d). Occasional crevasse-like features that had been labelled as supraglacial lakes (Figure 2c) were deleted, and bright bare ice mistakenly labelled as *superimposed ice* was reclassified as *ice* when the surrounding context (e.g. distance to *firn*) did not support refreezing (Figure 2e). At least two experts agreed on the assigned labels after the cross-check. Finally, additional points were added in areas where label density was low (Figure 2e,f). To quantify the consistency of the initial annotations, we report the expert agreement rate, defined as the percentage of labels that remained unchanged after the cross-check relative to the number of pre-cross-check labels. We also report the distribution shift, defined as the relative change in the number of labels per class after the cross-check. These statistics are summarised in Table 1. The high agreement rates (93.61% overall) indicate that the vast majority of labels were already consistent across experts. The lowest, yet still high, agreement was observed for *snow* (80.80%), stemming from many *firn* pixels being attributed to *snow* initially. Consistently, the largest distribution shift was observed for *firn* (+36.16%), reflecting both unified reinterpretation of previously ambiguous *snow–firn* boundaries and the addition of new points in regions with sparse labelling. The cross-check proved essential for maintaining consistent labelling, particularly given the subtle visual differences in some classes, e.g., *snow* vs. *firn* vs. *ice*.

A total of 138 273 labelled points were collected. The final dataset exhibits notable class imbalance, with *ice* constituting the largest proportion (37.11%), followed by *snow* (21.39%), *debris* (16.05%), *firn* (15.16%), *shadow* (7.18%), *superimposed ice* (1.70%), *cloud* (0.94%) and *water* (0.47%). Additionally, the spatial distribution of labels varies substantially across the 31 glaciers, reflecting their different surface areas and the number of acquisitions. As an extreme case, Skeidararjoekull accounts for the highest number of labels (29.44% of the total), while Maladeta features the smallest number of points (0.057%). The feature stack patches of 330 m × 330 m or 33 × 33 pixels around every point label were extracted as inputs for the glacier facies classification model.

The multi-glacier scope of this dataset and the variety of classes make it a valuable resource for algorithm benchmarking under heterogeneous conditions, bridging scales from local to multi-regional contexts. By combining multiple sensors, various acquisition dates and additional DEM-derived features, the dataset sets a robust foundation for further research in automated glacier mapping and related climate studies.

3. Methods

We adopted a straightforward, VGG-like CNN (Figure 3) to provide a baseline for automated classification of glacier surfaces (Simonyan and Zisserman, 2014). Developing the most advanced model was not the primary focus of this study. We, however, incorporated multiscale feature aggregation to preserve fine-grained information specific to the central pixel of each patch. The single central activations were sliced from the feature maps before the pooling layers and concatenated with the globally pooled feature vector. These shortcuts give the classifier access to both local detail and large-scale context, while intentionally breaking translation invariance, which is undesirable when the label pertains to a fixed, i.e. central, pixel location.

The network was optimised using Adam (Kingma and Ba, 2014) with a cosine decay learning rate schedule

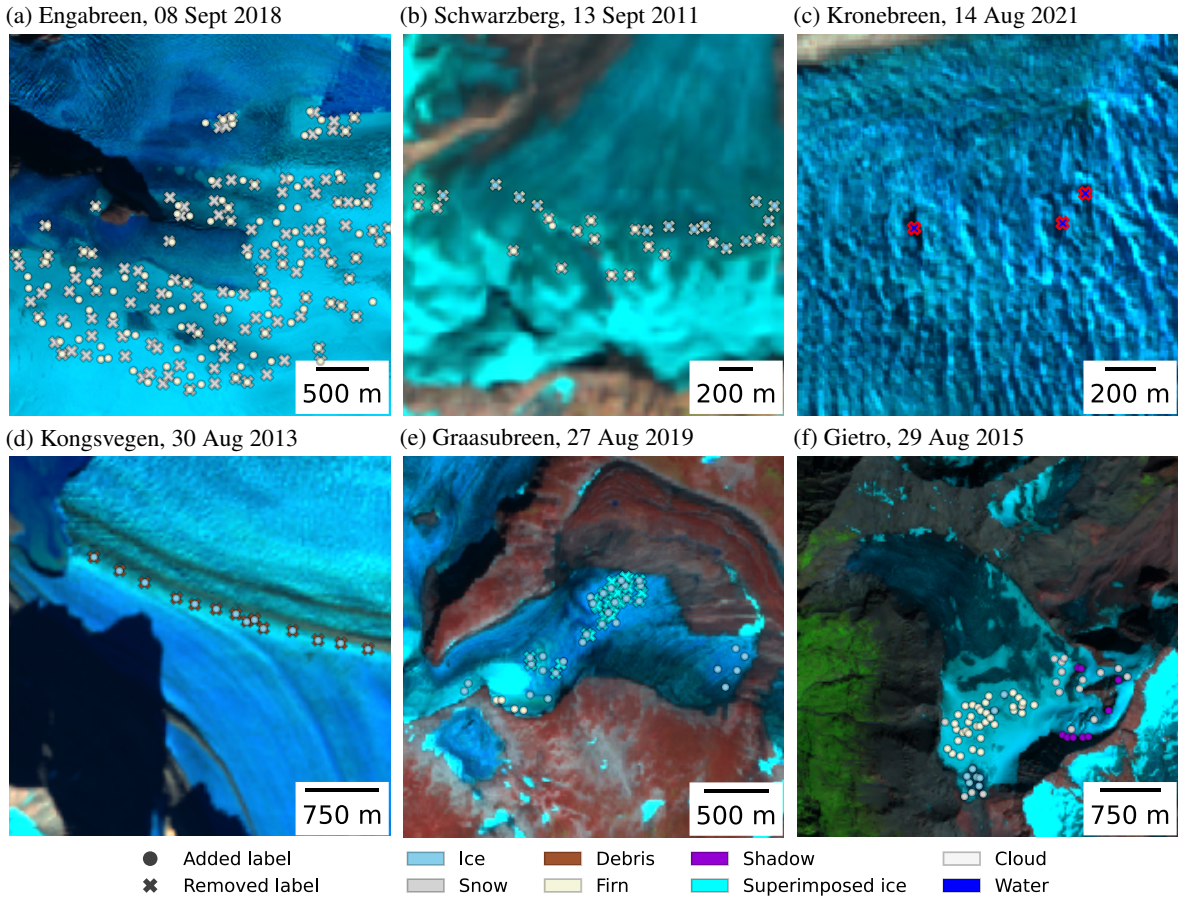

Figure 2: Examples of labels changed during the expert cross-check. The satellite images are presented in a false colour composition (R: $\text{SWIR}_{\approx 2.2\mu m}$, G: NIR, B: R). Landsat images courtesy of the U.S. Geological Survey. Copernicus Sentinel data 2015, 2018, 2019 and 2021.

(Loshchilov and Hutter, 2016) starting at $1e^{-3}$. Each training batch consisted of 512 samples. To improve model generalisation and mitigate overfitting, we applied data augmentation, including random brightness (± 0.01) and contrast adjustments ($\times 1 \pm 0.05$), horizontal/vertical flips and 90° rotations. This data augmentation is also applied to the elevation features and adds variability comparable to real thinning (Hugonnet et al., 2021), mitigating the issue of using a static DEM for mapping.

We trained the network by minimising a focal loss, a modification of cross-entropy that assigns higher weight to uncertain samples (Lin et al., 2017). In addition, we introduced weighting factors to account for the imbalance among the classes and the imbalance among different optical scenes. Let $c(i)$ be the class of sample i, and $s(i)$ the index of the optical scene from which sample i was taken. We define the weighting factors as $w_{c(i)} = N/N_{c(i)}$ and $w_{s(i)} = N/N_{s(i)}$, where N is the total number of samples in the training dataset, $N_{c(i)}$ is the number of samples belonging to class $c(i)$ and $N_{s(i)}$ is the number of samples extracted from scene $s(i)$. The overall loss function L is then defined as:

$$L_i = -\alpha\, w_{c(i)} w_{s(i)}\, \mathbf{y}_i \cdot \left[(1 - \hat{\mathbf{y}}_i)^{\gamma} \odot \log\left(\hat{\mathbf{y}}_i\right)\right] , \tag{1}$$

where $\alpha = 1$, $\gamma = 2$, $\odot$ stands for the Hadamard product, $\mathbf{y}_i$ are one-hot encoded true labels, and $\hat{\mathbf{y}}_i$ are model softmax outputs.

We split the 31 glaciers into six folds (Figure 1) to generate out-of-sample predictions for each glacier group. Within each iteration, one fold served as the out-of-sample test set, while the remaining five folds were subdivided into four folds for training and one for validation. The model was trained for 1000 epochs, and we selected the epoch with the highest Jaccard score averaged over classes on the validation set as the final model checkpoint. We repeated this procedure across all six folds to produce out-of-sample predictions for the entire dataset.

To enhance the reliability of the annotations, we applied a confident learning method (Northcutt et al., 2021) using the `cleanlab` package (Cleanlab, 2018). Confident learning identifies samples that are likely mislabelled by examining the confident joint matrix, which tabulates the joint distribution between predicted and given labels. Samples located in off-diagonal cells of the confident joint matrix are marked as confidently wrong and are pruned from the dataset to reduce potential label noise. In our study, we ran this procedure on the out-of-sample predictions from the nested cross-validation described above, pruning all samples falling into these off-diagonal cells except the samples of classes *shadow*, *superimposed ice*, *cloud* and *water*, which were assumed to be correctly labelled by the experts and relatively underrepresented, hence partly mitigating the class imbalance issue. This step ensures a higher degree of label consistency, thereby providing a cleaner benchmark.

We implemented the methodology from (Maslov et al., 2025) to report pixel-level confidence scores for the final predictions. The procedure corrects the raw self-confidence estimates so that they align with the observed validation accuracy. For every pixel of each validation fold, we recorded the raw confidence $s = \max_i \hat{\mathbf{y}}_i$ of the predicted class and whether the prediction was correct ($\text{acc} = \mathbb{1}\{\arg\max_i \hat{\mathbf{y}}_i = \arg\max_i \mathbf{y}_i\} \in \{0, 1\}$). We then fitted a kernel ridge regression model (Laplacian kernel with hyperparameters selected so that the models reconstructed the shapes of the curves almost identically) that learnt the mapping $s \mapsto \Pr(\text{acc} = 1 \mid s)$, i.e. predicted confidence to observed accuracy. The regressors were trained independently for each fold, and their output was clipped to $[0, 1]$ to ensure probabilistic interpretability. Applying these regressors to all pixels yielded calibrated confidence maps. We quantified

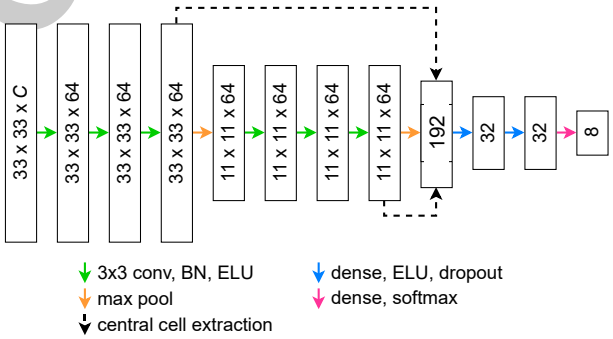

Figure 3: Convolutional neural network used in this study. Boxes represent tensors, arrows stand for operations, and numbers are the shapes of the tensors. C denotes the number of input features. ELU stands for the exponential linear unit (Clevert et al., 2015).

their reliability with expected calibration error. Unlike the original paper, where a Shannon-entropy-based confidence measure was used, we chose the maximum softmax score because it showed considerably lower expected calibration error even before any adjustment.

To provide an external consistency check and illustrate that our facies-derived variables capture glacier-wide SMB variability, we first calibrated a simple linear model similar to (Rabatel et al., 2005; Hock et al., 2007; Dumont et al., 2012) to reconstruct WGMS SMB records (WGMS, 2021) from purely remote sensing inputs:

$$\widehat{\mathrm{SMB}}_{t,i} = \beta_0 + \beta_1 \frac{A_{t,i}^{\mathrm{snow}}}{A_i} + \beta_2 \frac{h_{t,i}^{\mathrm{snowline}} - \overline{h}_i}{h_i^{\mathrm{max}} - h_i^{\mathrm{min}}} + \beta_3 \rho_{t,i} \,, \tag{2}$$

where $\beta_{0\text{–}3}$ are global model parameters shared by all glaciers, $A_{t,i}^{\mathrm{snow}}$ is the snow area, A_i is the glacier area, $h_{t,i}^{\mathrm{snowline}}$ is the transient snow line altitude computed as the 5-th percentile of the snow pixel elevations, $\overline{h}_i$ is the mean glacier elevation, h_i^{max} and h_i^{min} define the elevation range of the glacier, $\rho_{t,i}$ is the glacier-wide shortwave broadband albedo estimated as suggested by Liang (2001), t and i stand for year and glacier, respectively. Note that A_i and the hypsometric terms $\overline{h}_i$, h_i^{max}, h_i^{min} are kept constant; for each glacier i, we take the values from its earliest scene and use them for all subsequent years, so only the snow-dependent numerators vary with time. We adopted the six-fold cross-validation scheme described above—in every outer loop one glacier fold served as the hold-out test set, and the rest as training. This delivered out-of-sample SMB predictions for all glaciers without using any in-situ information from the test fold. Furthermore, because many earlier studies calibrated SMB lines on single glaciers, we also selected glaciers that have at least four annual WGMS records in our dataset (Kongsvegen, $n = 5$ from 1995, 2013, 2015, 2020 and 2021; Midtre Lovenbreen, $n = 4$ from 1995, 2013, 2015 and 2020) and fitted a reduced form:

$$\widehat{\mathrm{SMB}}_{t,i} = \beta_{0,i} + \beta_{1,i} \frac{A_{t,i}^{\mathrm{snow}}}{A_i} \,, \tag{3}$$

where $\beta_{0\text{–}1,i}$ capture climate and topography biases of each glacier. These models were evaluated in a leave-one-year-out manner—for every year, we trained on the remaining years of that glacier and predicted the left-out SMB. All SMB models were fitted with ridge regression, and the analysis was restricted to optical scenes acquired within a ±20-day window centred on 20 August, approximating the end of the ablation season.

4. Results

We evaluated the feature set configurations described above. For each feature set, we report classification performance both before (Table 2) and after (Table 3) applying the confident learning strategy to identify and prune confusing labels. Overall, model accuracy consistently improved in terms of the macro-average F_1 score in both cases when gradually adding more features, though class-specific metrics did not show consistent trends. The performance improved after confident learning, but it is important to note that the post-pruning results are not directly comparable to the pre-pruning results because the test folds were reduced (some potentially mislabelled or ambiguous samples were removed). To enable a fair comparison, we also evaluated the final model on the original (unpruned) test data

Table 2: Out-of-sample test performance before confident learning.

Feature set	Precision, %									Recall, %									F_1 score, %								
	ice	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average
1	88	**84**	84	55	60	53	**20**	49	62	81	**75**	67	66	96	**75**	35	24	65	85	79	74	60	74	62	**26**	33	62
2	90	81	84	55	92	47	14	29	62	82	62	90	**71**	96	66	**38**	25	66	**86**	70	87	62	94	55	21	27	63
3	**92**	82	83	**59**	90	42	15	39	63	81	70	**95**	70	96	73	33	26	**68**	86	76	89	**64**	93	54	21	31	64
4	83	82	**88**	58	**93**	**61**	13	**56**	**67**	**87**	67	77	65	**96**	74	35	**39**	67	85	74	82	61	**94**	**67**	19	**46**	**66**

Table 3: Out-of-sample test performance after confident learning. Note that feature sets are not directly comparable with each other, as well as with the results before confident learning, due to varying testing subsets. Classes *shadow*, *superimposed ice*, *cloud* and *water* were not pruned, and thus the corresponding pruning rates are not reported.

Feature set	Precision, %									Recall, %								F_1 score, %								Pruned, %				
	ice	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	Overall
1	94	91	95	79	87	60	16	50	71	90	88	88	85	94	77	35	26	73	92	89	91	82	90	68	22	34	71	12	21	5	20	13
2	94	92	92	86	85	59	30	53	74	90	91	90	89	96	77	37	17	73	92	92	91	87	90	67	33	26	72	11	31	4	21	15
3	93	93	94	79	92	86	44	50	79	94	86	96	84	96	72	37	23	73	94	89	95	81	94	78	40	32	75	12	25	2	21	13
4	95	93	94	80	95	69	54	80	82	93	90	96	88	96	72	35	41	76	94	92	95	84	96	70	42	54	78	9	29	14	24	16

(Table 4). Under this evaluation, feature set 4 also achieved the highest macro-average F_1 score at 69%, outperforming the other feature sets (65%, 64% and 67%, respectively) and the models before confident learning. This corresponds to an improvement of about 1–8% in F_1 score for each class compared to the model without confident learning, except for the *superimposed ice* class, which saw a drop of 4% in F_1 score. In other words, incorporating the full suite of features yielded the best overall classification performance.

Using feature set 4 after confident learning as the best one in general, we analysed the confusion matrix (Figure 4), which tabulates correct and misclassified samples for each class, to better understand class-wise performance. Notably, the confusion matrix before applying confident learning shows the same patterns as after confident learning. The most represented glacier facies—*ice* (F1 score = 94%), *snow* (92%), *debris* (95%), *firn* (84%) and *shadow* (96%)—were classified with high accuracy, as also indicated by high percentages along the matrix diagonal for these classes. However, the matrix also revealed systematic confusion between certain classes that aligns with their physical similarities. In particular, many pixels of *snow* were erroneously predicted as *firn* and vice versa. A similar overlap was

Table 4: Out-of-sample test performance after confident learning, evaluated on the unpruned dataset.

Feature set	Precision, %									Recall, %									F_1 score, %								
	ice	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average
1	90	86	88	62	85	49	12	39	64	81	**76**	87	73	94	77	35	26	69	85	**81**	87	67	89	60	18	31	65
2	91	86	85	**66**	84	49	12	38	64	83	73	88	**77**	96	**77**	37	17	68	87	79	86	**71**	90	60	18	23	64
3	**91**	86	86	63	91	**63**	14	43	67	85	70	**94**	74	96	72	**37**	23	69	**88**	77	90	68	93	**67**	20	30	67
4	90	**86**	**90**	62	**94**	56	**15**	**70**	**70**	**86**	71	91	76	**96**	72	35	**41**	**71**	88	78	**90**	68	**95**	63	21	**51**	**69**

Table 5: Per-region performance obtained with the best fold models.

Region	Precision, %									Recall, %									F_1 score, %									*n*								
	ice	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*sup.im.ice*	*cloud*	*water*	Overall
Pyrenees	98	0	–	25	–	–	–	–	41	100	0	–	50	–	–	–	–	50	99	0	–	33	–	–	–	–	44	55	3	0	2	0	0	0	0	60
Alps	97	96	94	86	93	17	68	0	69	96	87	94	95	94	19	86	0	71	97	91	94	90	94	18	76	0	70	10409	4219	1683	5078	1414	37	101	4	21389
Scandinavia	82	87	56	86	92	75	94	–	82	97	86	70	76	98	30	89	–	78	88	86	63	81	95	43	91	–	78	3778	2691	71	3294	420	517	18	0	10789
Iceland	96	93	95	82	85	43	32	88	77	94	90	97	90	88	83	10	65	77	95	92	96	86	87	57	15	75	75	23776	8330	16549	5465	1046	417	784	252	56619
Svalbard	95	96	86	58	97	84	60	79	82	85	94	92	86	98	86	69	25	79	90	95	89	70	97	85	64	38	78	8554	5644	750	2083	7046	1382	400	391	26250

observed between *firn* vs. *ice*, reflecting the gradual transition from older firn to bare glacier ice, as well as between *ice* and *superimposed ice*. These confusion trends are expected because glacier surface classes exhibit overlapping spectral characteristics and gradual boundaries. Likewise, a considerable portion of *water* was classified either as dark targets (*debris* or *shadow*) or as *superimposed ice*, as the latter closely resembles slush in satellite images. Another notable pattern was the misclassification of approximately two thirds of *cloud* pixels. Bruarjokull alone consists of 52.26% of all *cloud* labels in the dataset, most of which are very thin and semi-transparent allowing for the identification of the surface types beneath. The model attributed most of these pixels to the classes visible under the clouds (as seen in Figure 5j along the western boundary of the glacier). As a result, many true clouds were counted as *snow*, *ice*, *debris* or *firn*, driving down *cloud* recall and producing the main block of off-diagonal errors for that class. Excluding *cloud* pixels of the Bruarjokull case from the evaluation changes *cloud* precision from 54% to 50%, recall from 35% to 63% and F1 score from 42% to 56% (resulting in a macro-average F1 score of 80%). To assess the robustness of our method across contrasting glacierised landscapes, we evaluated the models separately for five region (Table 5). The models demonstrate a solid performance (F1 score $\geq$ 70%) for all cases with reasonable sample size ($n > 1000$), and less stable behaviour for the rest (0% $\leq$ F1 score $\leq$ 99%). Notably, the sample size (n) shows moderate correlation with F1 score ($r_s = 0.63$, $p \ll 0.001$). Glacier-wise macro-average F1 scores are 79.9 ± 11.5%.

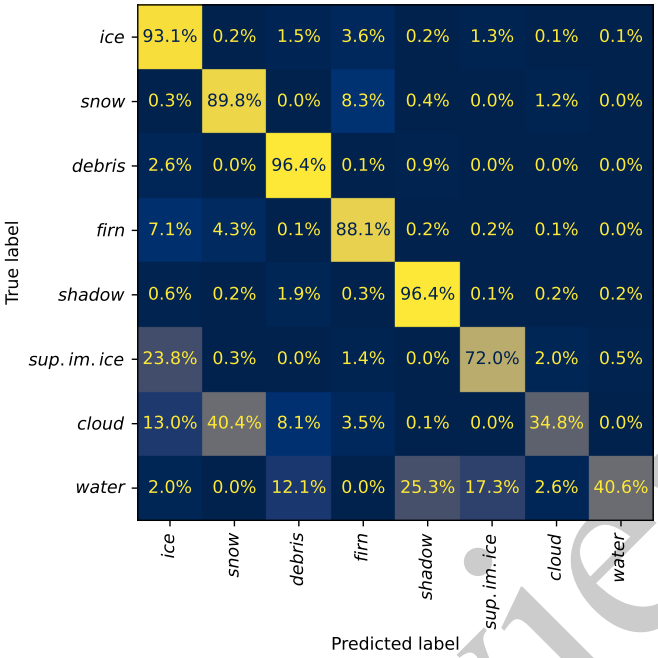

Figure 4: Confusion matrix for the best classification results obtained after confident learning, normalised per row.

We calibrated the model confidence scores to make them match real accuracies more closely. Before this step, the expected calibration error (ECE_{100}) was 11.5% (between 7%–13% and between 9%–23% across the test and validation folds, respectively). After the calibration procedure, the error fell to 3.6% (4–11% across test folds and 1–7% across validation folds). Figure 6 shows the corresponding reliability curves. Now, when the model assigns a label with an $X\%$ confidence to a pixel, that label is indeed correct in $\approx X\%$ cases. Well-calibrated scores are useful in practice because they highlight low-confidence areas that deserve closer inspection or extra labelling. The drop in calibration error confirms the trustworthiness of the calibrated predictive confidence. Calibration error, evaluated as the alignment between confidence and accuracy among samples predicted as a given class, remains satisfactory for *ice* (ECE_{100} = 3.7%), *snow* (6.0%), *debris* (6.5%), *shadow* (5.2%) and *water* (11%), but exceeds 20% for *firn* (21%), *superimposed ice* (27%) and *cloud* (62%; the latter being particularly high, aligned with the failure case discussed previously).

Qualitative evaluation of the classification maps (Figure 5) shows strong overall agreement with expert labels across diverse glacier scenes. The model correctly captured the broad spatial distribution of facies, delineating accumulation and ablation zones in accordance with manual reference maps. Higher predictive uncertainty was evident near class boundaries—for example, along the *snow-firn* and *firn-ice* interfaces—consistent with the inherently gradual transitions and hence higher expert disagreement in those areas (Figure 7a,b,c,d shows how the labels where experts had disagreements are generally located within or clustered around areas with low predictive confidence). Nevertheless, the majority of scenes were classified plausibly (Figure 5a-j). Two notable failure modes were identi-

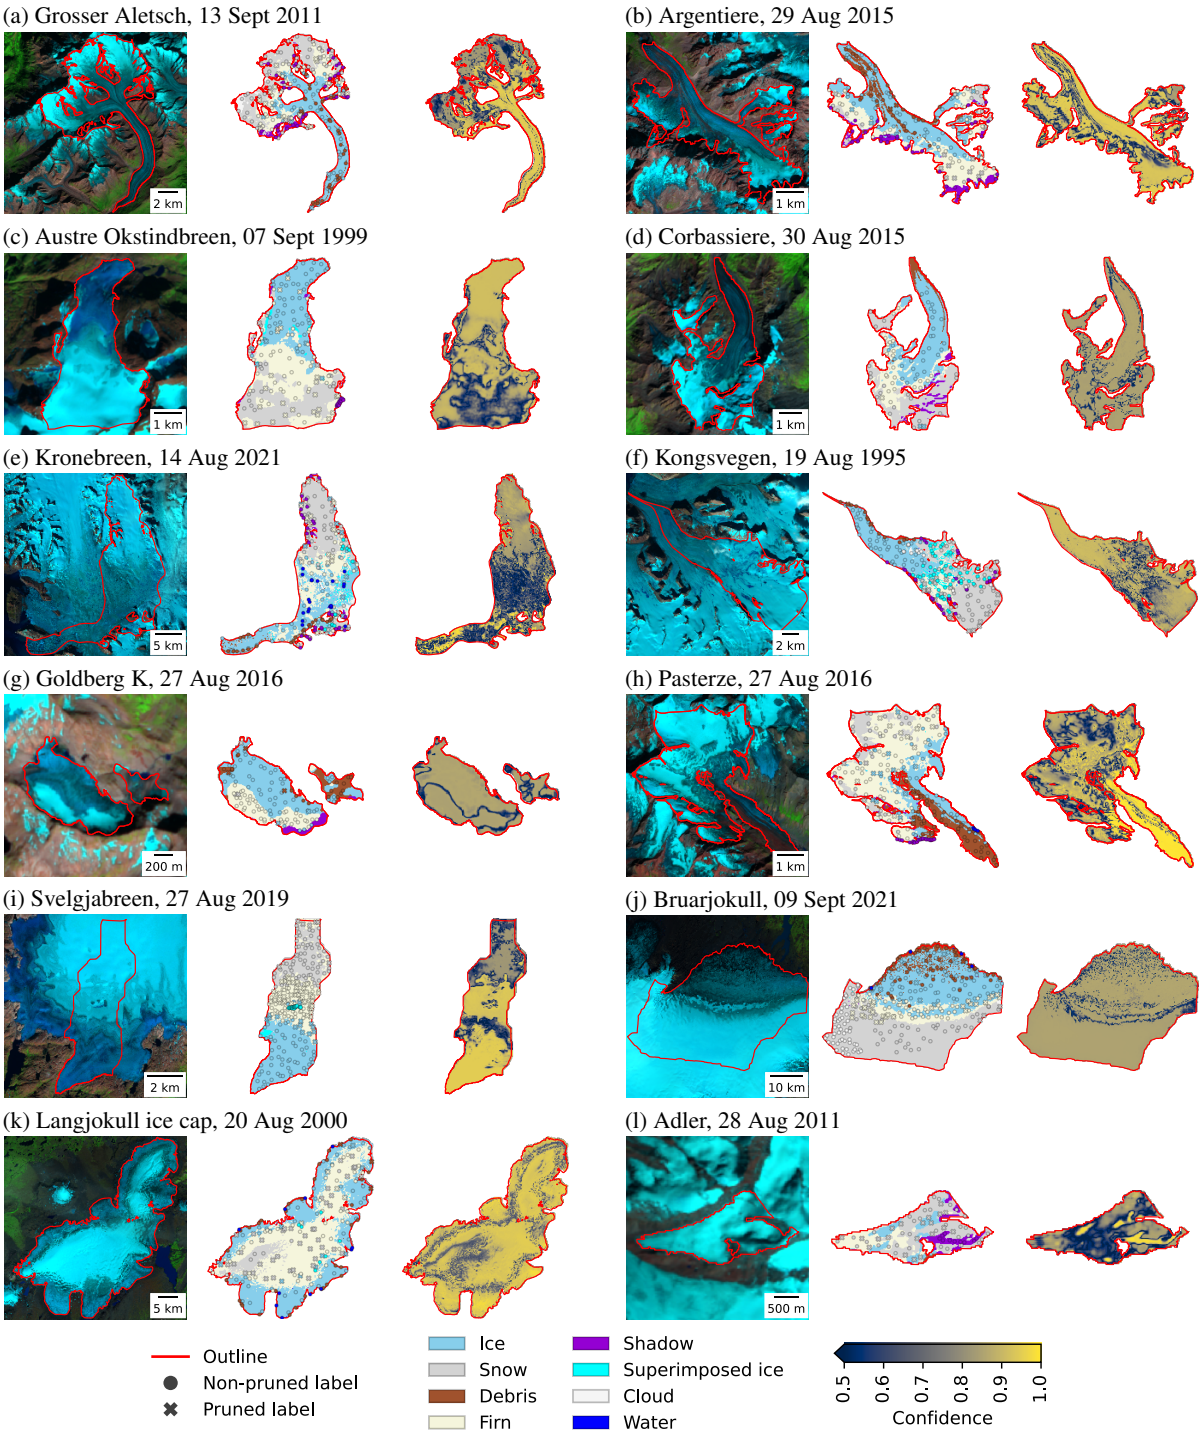

Figure 5: Qualitative view of the out-of-sample predictions for several glaciers. Each panel shows (left) an input satellite scene, (middle) predicted glacier facies maps and (right) predictive confidence. Markers in the middle images represent expert labels (a subset shown for clarity); their colour stands for the assigned class, and their shape indicates whether they were pruned via confident learning. The satellite images are presented in a false colour composition (R: $\text{SWIR}_{\approx 2.2\mu m}$, G: NIR, B: R). Landsat images courtesy of the U.S. Geological Survey. Copernicus Sentinel data 2015, 2016, 2019 and 2021.

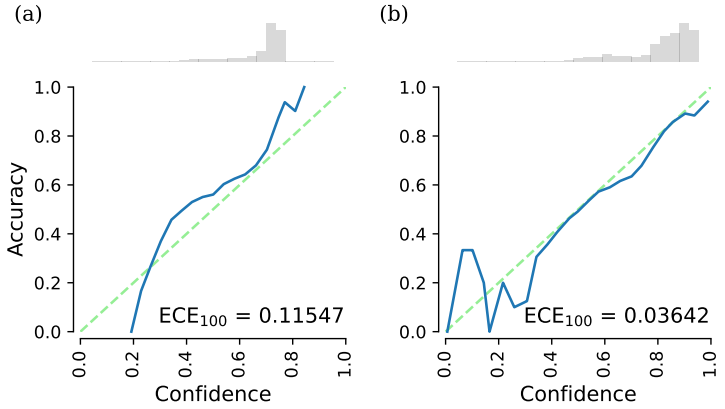

Figure 6: (a) Pre- and (b) post-calibration curves for the complete out-of-sample test data. The upper panels show the predictive confidence frequency. ECE_{100} stands for expected calibration error.

fied. In one case (the Langjokull 2000 scene, Figure 5k), the model facies predictions deviated substantially from the reference map attributing lots of *snow* pixels to *firn*; this anomaly may be due to an unusual radiometric appearance of the image, dust deposited on the snow surface or inconsistencies in the original labels. Notably, the majority of these labels were pruned after confident learning. In another scenario (the Adler 2011 scene, Figure 5l), scenes dominated by recent fresh snowfall caused the model to break down: the new snow cover made all surfaces look alike, and the classifier produced a seemingly random mix of *snow*, *firn* and *ice* labels. Aside from these outliers, the model demonstrated robust performance and generalised well to the varied conditions represented in the test folds.

Finally, we assessed the utility of the produced facies maps for SMB modelling. Figure 8 summarises the correspondence between facies-derived predictions and observed SMB. Using a single global model for all glaciers (Figure 8a, Equation 2), the facies data yielded a moderate, yet statistically significant Pearson correlation of $r = 0.63$ with measured SMB and an RMSE of 0.61 m w.e. When we calibrated a separate model for each glacier (Figure 8b, Equation 3), the correlation improved to $r = 0.81$ and RMSE = 0.26 m w.e. These values are on par with earlier studies that related facies to glacier mass balance with evaluation on a similar number of measurements per glacier (Kulkarni et al., 2004; Drolon et al., 2016), despite the differences in time ranges, methodology and validation protocols. The improved correlation for glacier-specific models is expected, since it accounts for the particular topographic and climatic characteristics of each glacier. Overall, this experiment demonstrates that our automatically generated facies maps can serve as informative inputs for SMB estimation. We stress, however, that this constitutes an illustrative validation and external consistency check; a full assessment of facies-based SMB modelling would require more extensive analysis beyond the scope of the current study.

5. Discussion & conclusions

This study delivers the first Europe-wide deep-learning solution for glacier facies mapping. We release both a comprehensive, multisensor dataset (the largest in terms of the number of glaciers, satellite scenes, classes and its

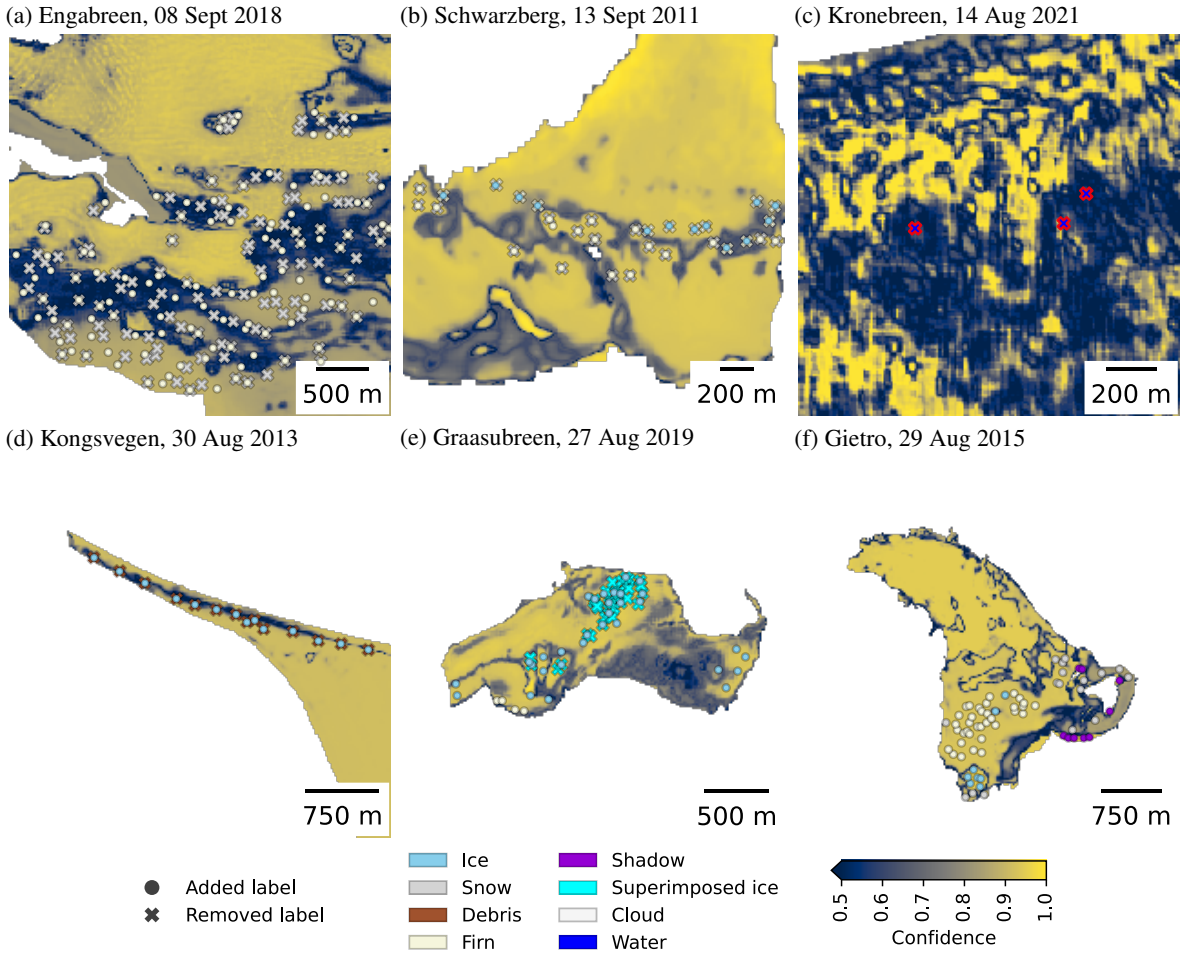

Figure 7: The same as in Figure 2, but shown against predictive confidence.

spatial extent) and a baseline model that already performs well across many glaciers, thanks to extra topographic features and label cleanup with confident learning. Our models learnt from dozens of scenes across different climates and years, demonstrating transferability and strong performance without site-specific tuning. The models yield an overall accuracy (98%) and macro-average F_1 score (78%) that are generally on par or surpass most previously reported results (Bhardwaj et al., 2015; Shukla and Yousuf, 2017; Aberle et al., 2025) Beyond headline scores, our dataset encompasses a wider set of classes, enabling more detailed interpretation of accumulation and ablation partitioning as well as refreezing intensity. We further propose a deliberately rigorous validation protocol by adopting nested cross-validation so every prediction is out-of-sample, better reflecting operational use, and report calibrated predictive uncertainty, which is rarely present in the facies literature, making our confidence maps actionable for mapping and quality control. Finally, we demonstrate external utility by reconstructing glacier-wide SMB from our classification products, indicating that our maps carry process-relevant signal at a broad scale. Researchers can run the models as-is or fine-tune them with modest local data, taking advantage of transfer learning. Together, the dataset and models form

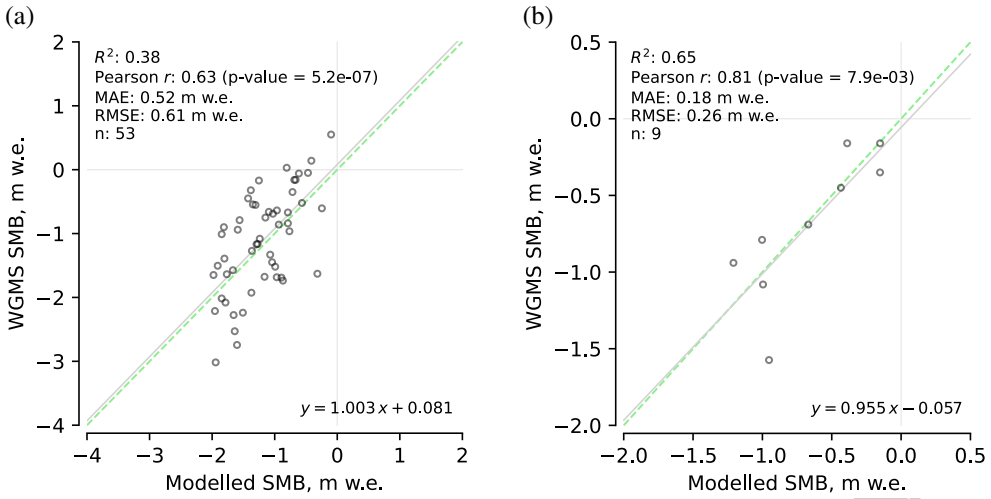

Figure 8: Glacier-wide SMB modelling results (a) of the six-fold cross-validation of the global model and (b) of the leave-one-year-out validation of the glacier-specific models for Kongsvegen and Midtre Lovenbreen.

a solid foundation for future glacier mapping work.

The regional analysis confirms that model robustness depends on data availability. The positive correlation between sample size and F1 score ($r_s = 0.63$, $p \ll 0.001$) indicates that regions with more extensive and diverse training data allow the model to better capture spectral and topographic variability, hence, expanding the dataset is a logical follow-up. Conversely, regions with limited glacierised area and fewer examples (e.g., the Pyrenees) show occasionally degraded results. Nonetheless, the glacier-wise macro-average F1 score of 79.9 ± 11.5% across all regions suggests that the method generalises well to distinct glacier settings from small temperate valley glaciers to large polythermal ice caps without site-specific tuning. The remaining regional spread likely reflects inherent differences in data availability and illumination geometry rather than systematic model bias, supporting the robustness and transferability of the methods across contrasting environments. In practice, one should expect even higher performance from ensembling the fold models in operational use.

Despite these positive results, several limitations are acknowledged. Glacier facies boundaries are inherently fuzzy in nature, which leads to ambiguity in the reference labels. For example, the transition between *snow* and *firn* is often gradual rather than abrupt, making it difficult even for experts to delineate a precise boundary. Consequently, some label noise and inter-expert disagreement are unavoidable in the reference labels, even given the performed cross-checks. Additionally, variations in imaging conditions and surface properties (e.g. wet snow during melt) can cause one facies to resemble another in the satellite data. These factors contribute to the observed confusion between *snow* vs. *firn* vs. *ice* classes, as also acknowledged by Rastner et al. (2019) and Aberle et al. (2025). The results from our three-expert, cross-checked protocol suggest that a larger, community-driven labelling effort with more experts could further improve uncertainty characterisation and support future extensions towards a global benchmark.

To address the challenges of label noise and class ambiguity, we adopted confident learning (Northcutt et al., 2021),

a data-centric artificial intelligence method. Rather than solely increasing model complexity, we focused on improving the quality of the training data and labels. In this study, confident learning identified and pruned mislabelled or highly ambiguous training samples. This process removed part of the problematic labels, which in turn improved the model performance and generalisation. These results show that investing effort into better labels and data consistency can improve model performance, complementing algorithmic advances. This reflection underscores a broader point: in domains like remote sensing and glaciology, where "groundtruth" is difficult and sometimes subjective, techniques to handle label uncertainty are essential.

Looking ahead, there are several avenues for improving glacier facies mapping. First, assembling a larger and more diverse training dataset would help capture the full variability of glacier conditions, potentially further boosting the performance and robustness of the model. Future work could expand the geographic scope and include more glaciers with different characteristics (including those with heavy debris cover and very steep slopes as observed in High-Mountain Asia, which we did not explicitly address in this study) and focus on adding more labels for heavily underrepresented classes, namely *superimposed ice*, *cloud* and *water*. Second, more advanced models should be explored. Likewise, complementary data such as synthetic-aperture radar imagery might provide signals to distinguish facies under challenging conditions, e.g. cloud cover or low-light months, with which optical imagery alone struggles. Improvements in auxiliary data could also enhance classification accuracy. A time-evolving DEM or glacier surface model would account for elevation changes over the years, alleviating errors introduced by our static DEM assumption. Additionally, post-processing steps could enforce spatial coherence in the output maps—for instance, by applying morphological filters to remove obvious outliers (e.g. small patches of *ice* predicted in the middle of a large *snow* field) or incorporating topological constraints (e.g. *snow* should generally lie above *firn* which lies above *ice* in elevation) to ensure that facies zones are contiguous and physically plausible. These enhancements, combined with our baseline, could further narrow the gap between automated classifications and expert interpretations.

The additional validation against independent SMB measurements shows that a simple linear model driven by features extracted from the CNN outputs (AAR, SLA) captures glacier-wide SMB variability. The glacier-specific models yield and RMSE of 0.26 m w.e. similar to other glacier-specific empirical methods (Kulkarni et al., 2004; Drolon et al., 2016), while the global model achieves an out-of-sample RMSE of 0.61 m w.e. This error is marginally larger than, e.g., the RMSE of 0.46 m w.e. reported by (Podsiadlo et al., 2020) for a much narrower latitudinal band and more complicated methodology involving physical modelling. Recent machine learning SMB reconstructions on point measurements confined to the European Alps driven by climatic and topographical predictors (Anilkumar et al., 2023; van der Meer et al., 2025) achieved comparable accuracies (RMSE $\approx$ 0.604–1.071 m w.e.), indicating common difficulties to generalise in multi-glacier settings. Our main limitations stem from the simplicity of the model, the use of a static DEM and from the sparse temporal sampling of late-season scenes. Denser facies time series and up-to-date elevation data are likely to yield significant accuracy gains. We see considerable potential in integrating the CNN-derived facies products and their estimated uncertainty into physics-informed SMB frameworks, enhancing transferability across diverse climatic regimes.

Overall, this study presents a new dataset of expertly labelled glacier facies across multiple European regions and a baseline deep learning model for automated facies classification. By means of feature design and data curation, including label refinement via confident learning, we achieved accurate facies maps for diverse glaciers (macro-average F_1 score = 78% after pruning). To our knowledge, this is one of the first studies to successfully apply deep learning at scale for glacier facies classification. We also illustrated how the resulting facies maps can be applied in a glaciological context—as proxies for assessing glacier surface mass balance with reasonably high fidelity. The methods and findings here lay the foundation for more advanced and large-scale facies mapping efforts, which will be crucial for climate change impact assessments and the monitoring of glacier health.

Data and code statement

The dataset compiled in this study is available at `https://doi.org/10.5281/zenodo.17473413`. Our codebase will be published upon acceptance of this paper at `https://github.com/konstantin-a-maslov/glacier_facies_classification`.

CRediT authorship contribution statement

Konstantin A. Maslov: Conceptualisation, Data curation, Methodology, Software, Writing—Original draft preparation. **Thomas Schellenberger:** Conceptualisation, Data curation, Supervision, Writing—Review & editing, Funding acquisition. **Prashant Pandit:** Data curation, Writing—Review & editing. **Claudio Persello:** Conceptualisation, Supervision, Writing—Review & editing, Funding acquisition. **Alfred Stein:** Supervision, Writing—Review & editing.

Declaration of competing interest

The authors declare no conflict of interest.

Funding sources

This research was financed by the Research Council of Norway under the "Researcher Project for Scientific Renewal" (project MASSIVE, no. 315971). Additional support was provided by Open Clouds for Research Environments under the EU H2020 programme (project MATS_CLOUD, no. 824079) for access to the CREODIAS cloud computing platform.

References

Aberle, R., Enderlin, E., O'Neel, S., Florentine, C., Sass, L., Dickson, A., Marshall, H.P., Flores, A., 2025. Automated snow cover detection on mountain glaciers using spaceborne imagery and machine learning. The Cryosphere 19, 1675–1693. URL: https://tc.copernicus.org/articles/19/1675/2025/, doi:10.5194/tc-19-1675-2025.

Ali, I., Shukla, A., Romshoo, S.A., 2017. Assessing linkages between spatial facies changes and dimensional variations of glaciers in the upper indus basin, western himalaya. Geomorphology 284, 115–129. URL: https://www.sciencedirect.com/science/article/pii/S0169555X17300132, doi:https://doi.org/10.1016/j.geomorph.2017.01.005.

Anilkumar, R., Bharti, R., Chutia, D., Aggarwal, S.P., 2023. Modelling point mass balance for the glaciers of the central european alps using machine learning techniques. The Cryosphere 17, 2811–2828. URL: https://tc.copernicus.org/articles/17/2811/2023/, doi:10.5194/tc-17-2811-2023.

Bhardwaj, A., Joshi, P., Snehmani, Sam, L., Singh, M.K., Singh, S., Kumar, R., 2015. Applicability of landsat 8 data for characterizing glacier facies and supraglacial debris. International Journal of Applied Earth Observation and Geoinformation 38, 51–64. URL: https://www.sciencedirect.com/science/article/pii/S0303243414002864, doi:https://doi.org/10.1016/j.jag.2014.12.011.

Carrivick, J.L., Tweed, F.S., 2016. A global assessment of the societal impacts of glacier outburst floods. Global and Planetary Change 144, 1–16. URL: https://www.sciencedirect.com/science/article/pii/S0921818116301023, doi:https://doi.org/10.1016/j.gloplacha.2016.07.001.

Cleanlab, 2018. cleanlab. https://github.com/cleanlab/cleanlab.

Clevert, D.A., Unterthiner, T., Hochreiter, S., 2015. Fast and accurate deep network learning by exponential linear units (elus), in: 4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings. URL: https://arxiv.org/abs/1511.07289v5.

Dietz, A.J., Kuenzer, C., Gessner, U., Dech, S., 2011. Remote sensing of snow—a review of available methods. International Journal of Remote Sensing 33, 4094–4134. URL: https://doi.org/10.1080/01431161.2011.640964, doi:10.1080/01431161.2011.640964.

Drolon, V., Maisongrande, P., Berthier, E., Swinnen, E., Huss, M., 2016. Monitoring of seasonal glacier mass balance over the european alps using low-resolution optical satellite images. Journal of Glaciology 62, 912–927. URL: https://www.cambridge.org/core/journals/journal-of-glaciology/article/monitoring-of-seasonal-glacier-mass-balance-over-the-european-alps-using-lowresolution-optical-sate
2AD9341322E5B40A9FDA03A6CB6B2E5E, doi:10.1017/JOG.2016.78.

Dumont, M., Gardelle, J., Sirguey, P., Guillot, A., Six, D., Rabatel, A., Arnaud, Y., 2012. Linking glacier annual mass balance and glacier albedo retrieved from modis data. Cryosphere 6, 1527–1539. doi:10.5194/TC-6-1527-2012.

Florath, J., Keller, S., Staub, G., Weinmann, M., 2021. Optical remote sensing for glacier monitoring with respect to different snow and ice types: A case study for the southern patagonian icefield, in: 2021 11th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS), pp. 1–5. doi:10.1109/WHISPERS52202.2021.9484055.

Fox, E., Schwartz-Marin, E., Rangecroft, S., Palmer, S., Harrison, S., 2024. 'water resource' framing for the value and governance of glacier water availability in the semi-arid chilean andes. Frontiers in Water Volume 6. URL: https://www.frontiersin.org/journals/water/articles/10.3389/frwa.2024.1367889, doi:10.3389/frwa.2024.1367889.

GLIMS Consortium, 2015. GLIMS Glacier Database, Version 1. Boulder Colorado, USA. NASA National Snow and Ice Data Center Distributed Active Archive Center. doi:10.7265/N5V98602.

Gupta, R., Haritashya, U., Singh, P., 2005. Mapping dry/wet snow cover in the indian himalayas using irs multispectral imagery. Remote Sensing of Environment 97, 458–469. URL: https://www.sciencedirect.com/science/article/pii/S0034425705001574, doi:https://doi.org/10.1016/j.rse.2005.05.010.

Hall, D., Ormsby, J., Bindschadler, R., Siddalingaiah, H., 1987. Characterization of snow and ice reflectance zones on glaciers using landsat thematic mapper data. Annals of Glaciology 9, 104–108. doi:10.3189/S0260305500000471.

Hall, D.K., Chang, A.T., Siddalingaiah, H., 1988. Reflectances of glaciers as calculated using landsat-5 thematic mapper data. Remote Sensing of Environment 25, 311–321. URL: https://www.sciencedirect.com/science/article/pii/0034425788901071, doi:https://doi.org/10.1016/0034-4257(88)90107-1.

Herreid, S., Pellicciotti, F., 2020. The state of rock debris covering earth's glaciers. Nature Geoscience 13, 621–627. URL: https://www.nature.com/articles/s41561-020-0615-0, doi:10.1038/s41561-020-0615-0.

Hock, R., Kootstra, D.S., Reijmer, C., 2007. Deriving glacier mass balance from accumulation area ratio on storglaciären, sweden, in: 7th Scientific Assembly of the International Association of Hydrological Science, IAHS - Workshop on Andean Glaciology and Symposium on the Contribution from Glaciers and Snow Cover to Runoff from Mountains in Different Climates, pp. 163–170. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-38549156622&partnerID=40&md5=abe3ce04fbedff2850dd16badab023de.

Hugonnet, R., McNabb, R., Berthier, E., Menounos, B., Nuth, C., Girod, L., Farinotti, D., Huss, M., Dussaillant, I., Brun, F., Kääb, A., 2021. Accelerated global glacier mass loss in the early twenty-first century. Nature 592, 726–731. doi:10.1038/S41586-021-03436-Z.

Huo, D., Bishop, M.P., Bush, A.B.G., 2021. Understanding complex debris-covered glaciers: Concepts, issues, and research directions. Frontiers in Earth Science Volume 9. URL: https://www.frontiersin.org/journals/earth-science/articles/10.3389/feart.2021.652279, doi:10.3389/feart.2021.652279.

Jawak, S.D., Wankhede, S.F., Luis, A.J., Balakrishna, K., 2022. Multispectral characteristics of glacier surface facies (chandra-bhaga basin, himalaya, and ny-Ålesund, svalbard) through investigations of pixel and object-based mapping using variable processing routines. Remote Sensing 14. URL: https://www.mdpi.com/2072-4292/14/24/6311, doi:10.3390/rs14246311.

Kingma, D.P., Ba, J.L., 2014. Adam: A method for stochastic optimization, in: 3rd International Conference on Learning Representations, ICLR 2015—Conference Track Proceedings, International Conference on Learning Representations, ICLR. URL: https://arxiv.org/abs/1412.6980v9, doi:10.48550/arxiv.1412.6980.

Kulkarni, A.V., Rathore, B.P., Alex, S., 2004. Monitoring of glacial mass balance in the baspa basin using accumulation area ratio method. Current Science 86, 185–190. URL: http://www.jstor.org/stable/24109532.

König, M., Winther, J.G., Kohler, J., König, F., 2004. Two methods for firn-area and mass-balance monitoring of svalbard glaciers with sar satellite images. Journal of Glaciology 50, 116–128. doi:10.3189/172756504781830286.

Larocca, L.J., Lea, J.M., Erb, M.P., McKay, N.P., Phillips, M., Lamantia, K.A., Kaufman, D.S., 2024. Arctic glacier snowline altitudes rise 150 m over the last 4 decades. The Cryosphere 18, 3591–3611. URL: https://tc.copernicus.org/articles/18/3591/2024/, doi:10.5194/tc-18-3591-2024.

Liang, S., 2001. Narrowband to broadband conversions of land surface albedo i: Algorithms. Remote Sensing of Environment 76, 213–238. URL: https://www.sciencedirect.com/science/article/pii/S0034425700002054, doi:https://doi.org/10.1016/S0034-4257(00)00205-4.

Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2017. Focal loss for dense object detection, in: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 2999–3007. doi:10.1109/ICCV.2017.324.

Loibl, D., Richter, N., Grünberg, I., 2025. Remote sensing-derived time series of transient glacier snowline altitudes for high mountain asia, 1985–2021. Scientific Data 12, 103. URL: https://doi.org/10.1038/s41597-024-04309-6, doi:10.1038/s41597-024-04309-6.

Loshchilov, I., Hutter, F., 2016. Sgdr: Stochastic gradient descent with warm restarts, in: 5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings, International Conference on Learning Representations, ICLR. URL: https://arxiv.org/abs/1608.03983v5.

Lu, Y., Zhang, Z., Shangguan, D., Yang, J., 2021. Novel machine learning method integrating ensemble learning and deep learning for mapping debris-covered glaciers. Remote Sensing 13, 2595. URL: https://www.mdpi.com/2072-4292/13/13/2595, doi:10.3390/RS13132595.

Luis, A., Singh, S., 2020. High-resolution multispectral mapping facies on glacier surface in the arctic using worldview-3 data. Czech Polar Reports 10, 23–26. doi:10.5817/CPR2020-1-3.

Maslov, K.A., Persello, C., Schellenberger, T., Stein, A., 2025. Globally scalable glacier mapping by deep learning matches expert delineation accuracy. Nature Communications 16, 43. URL: https://doi.org/10.1038/s41467-024-54956-x, doi:10.1038/s41467-024-54956-x.

van der Meer, M., Zekollari, H., Huss, M., Bolibar, J., Sjursen, K.H., Farinotti, D., 2025. A minimal machine-learning glacier mass balance model. The Cryosphere 19, 805–826. URL: https://tc.copernicus.org/articles/19/805/2025/, doi:10.5194/tc-19-805-2025.

Müller, F., 1962. Zonation in the accumulation area of the glaciers of axel heiberg island, n.w.t., canada. Journal of Glaciology 4, 302–311. doi:10.3189/S0022143000027623.

Northcutt, C.G., Jiang, L., Chuang, I.L., 2021. Confident learning: Estimating uncertainty in dataset labels. Journal of Artificial Intelligence Research 70, 1373–1411. URL: https://jair.org/index.php/jair/article/view/12125, doi:10.1613/JAIR.1.12125.

Pfau, M., Veh, G., Schwanghart, W., 2023. Cast shadows reveal changes in glacier surface elevation. The Cryosphere 17, 3535–3551. URL: https://tc.copernicus.org/articles/17/3535/2023/, doi:10.5194/tc-17-3535-2023.

Podsiadlo, I., Paris, C., Callegari, M., Marin, C., Günther, D., Strasser, U., Notarnicola, C., Bruzzone, L., 2020. Integrating models and remote sensing data for distributed glacier mass balance estimation. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 13, 6177–6194. doi:10.1109/JSTARS.2020.3028653.

Pope, A., Rees, G., 2014. Using in situ spectra to explore landsat classification of glacier surfaces. International Journal of Applied Earth Observation and Geoinformation 27, 42–52. URL: https://www.sciencedirect.com/science/article/pii/S0303243413000937, doi:https://doi.org/10.1016/j.jag.2013.08.007. special Issue on Polar Remote Sensing 2013.

Prieur, C., Rabatel, A., Thomas, J.B., Farup, I., Chanussot, J., 2022. Machine learning approaches to automatically detect glacier snow lines on multi-spectral satellite images. Remote Sensing 14. URL: https://www.mdpi.com/2072-4292/14/16/3868, doi:10.3390/rs14163868.

Rabatel, A., Dedieu, J.P., Vincent, C., 2005. Using remote-sensing data to determine equilibrium-line altitude and mass-balance time series: validation on three french glaciers, 1994–2002. Journal of Glaciology 51, 539–546. URL: https://www.cambridge.org/core/journals/journal-of-glaciology/article/using-remotesensing-data-to-determine-equilibriumline-altitude-and-massbalance-time-series-validati7EE925C092A3DB1E15EBDAE02E769F68, doi:10.3189/172756505781829106.

Rabatel, A., Sirguey, P., Drolon, V., Maisongrande, P., Arnaud, Y., Berthier, E., Davaze, L., Dedieu, J.P., Dumont, M., 2017. Annual and seasonal glacier-wide surface mass balance quantified from changes in glacier surface state: A review on existing methods using optical satellite imagery. Remote Sensing 9. doi:10.3390/rs9050507.

Rastner, P., Prinz, R., Notarnicola, C., Nicholson, L., Sailer, R., Schwaizer, G., Paul, F., 2019. On the automated mapping of snow cover on glaciers and calculation of snow line altitudes from multi-temporal landsat data. Remote Sensing 11. URL: https://www.mdpi.com/2072-4292/11/12/1410, doi:10.3390/rs11121410.

Rounce, D.R., Hock, R., Maussion, F., Hugonnet, R., Kochtitzky, W., Huss, M., Berthier, E., Brinkerhoff, D., Compagno, L., Copland, ., Farinotti, D., B., McNabb, R.W., 2023. Global glacier change in the 21st century: Every increase in temperature matters. Science 379, 78–83. URL: https://www.science.org/doi/abs/10.1126/science.abo1324, doi:10.1126/science.abo1324.

de Ruijter, T., 2016. python-dem-shadows. https://github.com/tomderuijter/python-dem-shadows.

Shukla, A., Yousuf, B., 2017. Evaluation of multisource data for glacier terrain mapping: a neural net approach. Geocarto International 32, 569–587. URL: https://doi.org/10.1080/10106049.2016.1161078, doi:10.1080/10106049.2016.1161078.

Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition, in: 3rd International Conference on Learning Representations, ICLR 2015—Conference Track Proceedings, International Conference on Learning Representations, ICLR. URL: https://arxiv.org/abs/1409.1556, doi:10.48550/arXiv.1409.1556.

Wendleder, A., Schmitt, A., Erbertseder, T., D'Angelo, P., Mayer, C., Braun, M.H., 2021. Seasonal evolution of supraglacial lakes on baltoro glacier from 2016 to 2020. Frontiers in Earth Science Volume 9 - 2021. URL: https://www.frontiersin.org/journals/earth-science/articles/10.3389/feart.2021.725394, doi:10.3389/feart.2021.725394.

WGMS, 2021. Fluctuations of Glaciers Database. World Glacier Monitoring Service, Zurich, Switzerland. doi:10.5904/wgms-fog-2021-05.

Williams, R.S., Hall, D.K., Benson, C.S., 1991. Analysis of glacier facies using satellite techniques. Journal of Glaciology 37, 120–128. doi:10.3189/S0022143000042878.

Wright, A., Wadham, J., Siegert, M., Luckman, A., Kohler, J., 2005. Modelling the impact of superimposed ice on the mass balance of an arctic glacier under scenarios of future climate change. Annals of Glaciology 42, 277–283. doi:10.3189/172756405781813104.

Yousuf, B., Shukla, A., Arora, M.K., Jasrotia, A.S., 2019. Glacier facies characterization using optical satellite data: Impacts of radiometric resolution, seasonality, and surface morphology. Progress in Physical Geography:

Earth and Environment 43, 473–495. URL: https://doi.org/10.1177/0309133319840770, doi:10.1177/0309133319840770.

Zemp, M., Huss, M., Thibert, E., Eckert, N., McNabb, R., Huber, J., Barandun, M., Machguth, H., Nussbaumer, S.U., Gärtner-Roer, I., Thomson, L., Paul, F., Maussion, F., Kutuzov, S., Cogley, J.G., 2019. Global glacier mass changes and their contributions to sea-level rise from 1961 to 2016. Nature 568, 382–386. doi:10.1038/S41586-019-1071-0.

Zemp, M., Jakob, L., Dussaillant, I., Nussbaumer, S.U., Gourmelen, N., Dubber, S., A, G., Abdullahi, S., Andreassen, L.M., Berthier, E., Bhattacharya, A., Blazquez, A., Boehm Vock, L.F., Bolch, T., Box, J., Braun, M.H., Brun, F., Cicero, E., Colgan, W., Eckert, N., Farinotti, D., Florentine, C., Floricioiu, D., Gardner, A., Harig, C., Hassan, J., Hugonnet, R., Huss, M., Jóhannesson, T., Liang, C.C.A., Ke, C.Q., Khan, S.A., King, O., Kneib, M., Krieger, L., Maussion, F., Mattea, E., McNabb, R., Menounos, B., Miles, E., Moholdt, G., Nilsson, J., Pálsson, F., Pfeffer, J., Piermattei, L., Plummer, S., Richter, A., Sasgen, I., Schuster, L., Seehaus, T., Shen, X., Sommer, C., Sutterley, T., Treichler, D., Velicogna, I., Wouters, B., Zekollari, H., Zheng, W., Team, T., 2025. Community estimate of global glacier mass changes from 2000 to 2023. Nature 639, 382–388. URL: https://doi.org/10.1038/s41586-024-08545-z, doi:10.1038/s41586-024-08545-z.

Zhang, J., Jia, L., Menenti, M., Hu, G., 2019. Glacier facies mapping using a machine-learning algorithm: The parlung zangbo basin case study. Remote Sensing 11. URL: https://www.mdpi.com/2072-4292/11/4/452, doi:10.3390/rs11040452.