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

Konstantin A. Maslov[1], Thomas Schellenberger[2], Prashant Pandit[3,4], Claudio Persello[1], and Alfred Stein[1]

[1]Department of Earth Observation Science, Faculty of Geo-information Science and Earth Observation (ITC), University of Twente, Overijssel, The Netherlands
[2]Department of Geosciences, Faculty of Mathematics and Natural Sciences, University of Oslo, Østlandet, Norway
[3]Department of Information Engineering and Computer Science, University of Trento, Trento, Trentino, Italy
[4]Institute for Earth Observation, Eurac Research, Bolzano, South Tyrol, Italy

Correspondence: Konstantin A. Maslov (k.a.maslov@utwente.nl)

Abstract. Glacier facies play a critical role in understanding the mass balance of glaciers, offering insights into accumulation and melting processes. Large-scale mapping of glacier facies from satellite data 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, 137 592 expert point labels and eight classes—five glacier facies (*ice*, *snow*, *debris*, *firn* and *refrozen-like*) and three miscellaneous classes (*shadow*, *water* and *cloud*)—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 82% on the complete cleaned data or 74% on the full, unpruned data, and $82.3 \pm 10.5\%$ glacier-wise. This performance remains consistent across different regions and sensors. 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 ($r = 0.65$, $\text{RMSE} = 0.60$ m w.e., where 1 m w.e. $= 1000$ kg m^{-2} denotes metre water equivalent) and competitive correspondence for glacier-specific calibration ($r = 0.79$, $\text{RMSE} = 0.28$ 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.

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 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.

25   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

30   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

35   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) by means of 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;

40   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

45   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.

50   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

55   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 delivers 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 (Fig. 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, 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 across 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 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
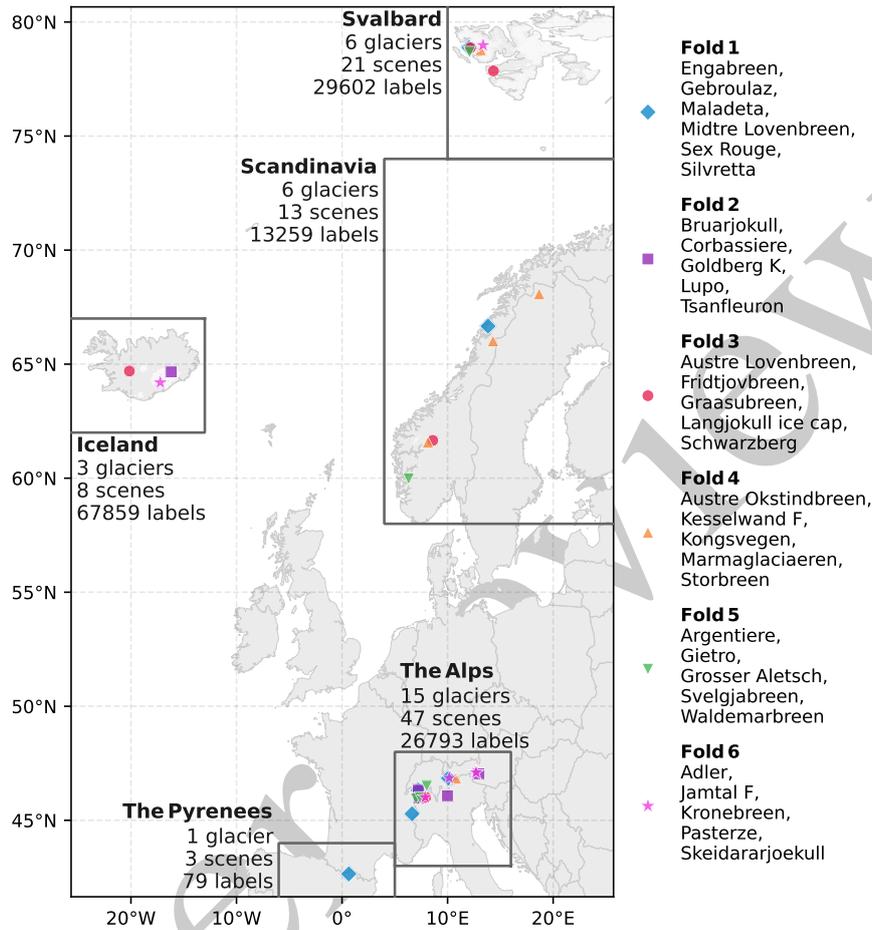
3

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

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 visually inspected 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

All images were resampled to $10\,\text{m} \times 10\,\text{m}$ grids with bilinear interpolation to match the spatial resolution of the Sentinel-2 data.

4

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

	ice	*snow*	*debris*	*firn*	*shadow*	*refrozen*	*water*	*cloud*	Total
Expert agreement, %	99.96	80.80	99.04	91.76	99.96	96.22	84.26	100.00	93.64
Distribution shift, %	+3.92	-19.05	-0.61	+36.16	+0.11	-0.13	-14.42	+10.48	+0.27

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 *refrozen-like* ice. *Ice* refers to exposed bare glacier ice free of snow cover. *Snow* is seasonal accumulation from the current hydrological year. *Firn* is older snow that has survived at least one ablation season and begun densification. *Refrozen-like* ice denotes refrozen meltwater forming a transparent ice layer on ice or firn and is an indicator of superimposed ice. *Debris* represents supraglacial rock or sediment cover. The remaining three classes are not glacier facies: *shadow* (areas affected by topographic shading), *water* (supraglacial and proglacial lakes) and *cloud* (any identifiable cloud cover).

Labelling was conducted on each optical scene, focusing on the glacier surfaces. These class assignments reflect expert visual interpretation of surface conditions in the optical imagery and thus constitute observational proxies that may not perfectly correspond to physically defined glacier facies in all cases. Typical spectral thresholds served as preliminary guidance for the main classes: pixels with NIR reflectance $\gtrsim 0.55$ were labelled *snow*, $\approx 0.40\ldots0.55$ as *firn*, and $\lesssim 0.40$ 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. *Refrozen-like* was recognised by its distinctive bright-blue appearance in the visible part of the spectrum (R-G-B) and deep-blue shades in SWIR-NIR-R composites, consistent with Winther (1993). Additional cues included yellow-brown appearance for *firn*, pronounced crevasse texture for *ice*, reddish SWIR-NIR-R 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 *refrozen-like* 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 (e.g. messy *firn–ice* transitions or very thin transparent clouds), 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* (Fig. 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* (Fig. 2d). Occasional crevasse-like features that had been labelled as supraglacial lakes (Fig. 2c) were deleted, and bright bare ice mis-
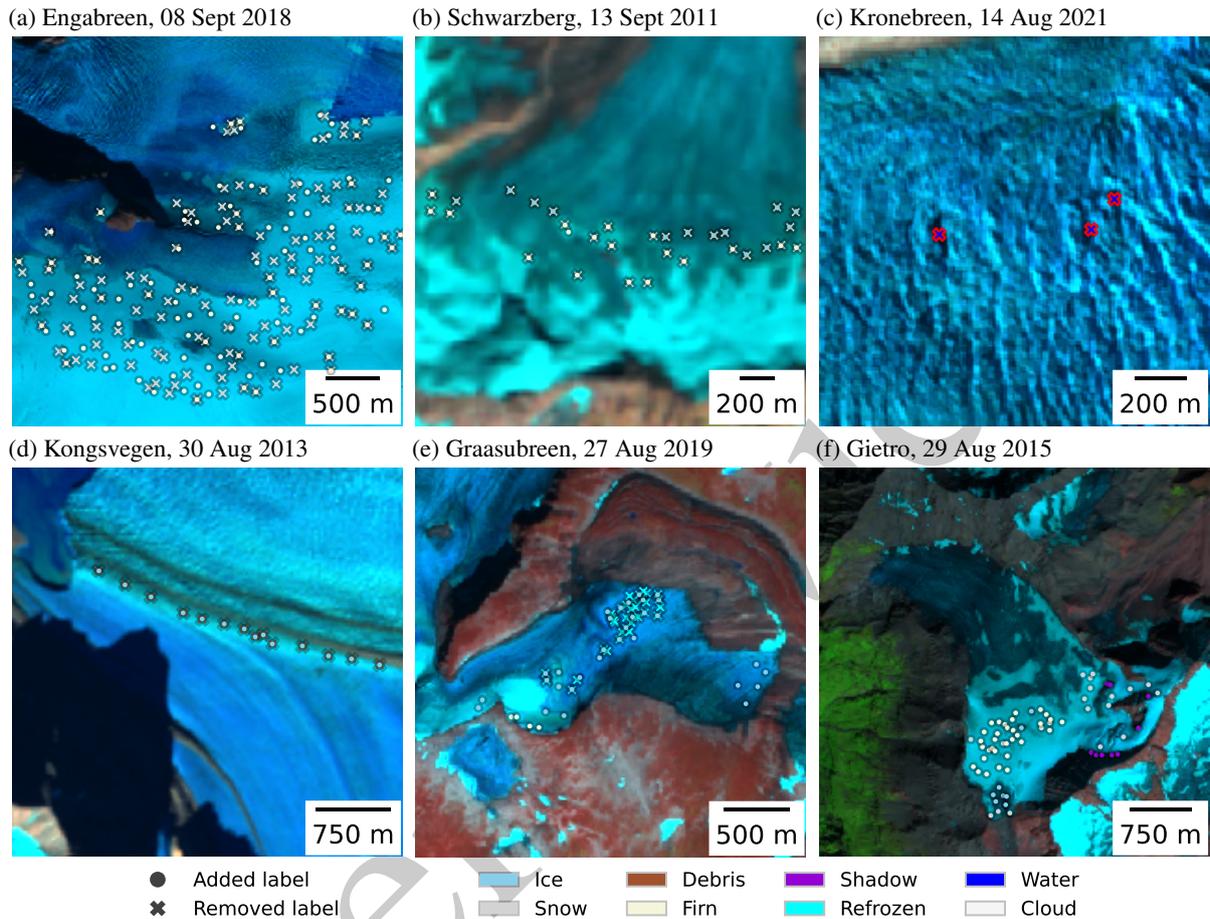
Figure 2. Examples of labels changed during the expert cross-check. The satellite images are presented in a false colour composition (R: SWIR$_{\approx 2.2 \mu m}$, G: NIR, B: R). Landsat images courtesy of the U.S. Geological Survey. Copernicus Sentinel data 2015, 2018, 2019 and 2021.

takenly labelled as *refrozen-like* was reclassified as *ice* when the surrounding context (e.g. distance to *firn*) did not support refreezing (Fig. 2e). At least two experts agreed on the assigned labels after the cross-check. Finally, additional points were

130   added in areas where label density was low (Fig. 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 (93.64% overall) indicate that the vast majority of labels were already consistent across experts. The lowest, yet still high, agreement was observed for *snow*

135   (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

6

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 137 592 labelled points were collected. The final dataset exhibits notable class imbalance, with *ice* constituting the largest proportion (37.29%), followed by *snow* (21.50%), *debris* (16.13%), *firn* (15.23%), *shadow* (7.21%), *refrozen-like* (1.71%), *water* (0.47%) and *cloud* (0.45%). 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 (29.58% of the total), while Maladeta features the smallest number of points (0.057%). The feature stack patches of $330\,\text{m} \times 330\,\text{m}$ 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 (Fig. 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 finegrained 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 largescale 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 (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 ($\times 1\pm0.05$), horizontal/vertical flips and $90°$ rotations. This data augmentation was 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:

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

where $\alpha = 1$, $\gamma = 2$, $\odot$ stands for the Hadamard product, $\mathbf{y}_i$ are one-hot encoded true labels, and $\widehat{\mathbf{y}}_i$ are model softmax outputs.
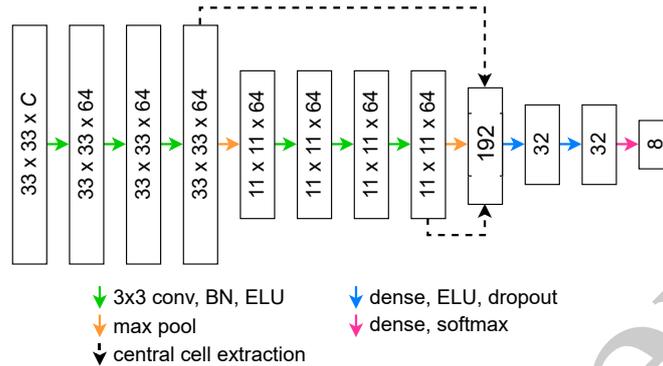
Figure 3. Convolutional neural network used in this study. Boxes represent tensors, arrows stand for operations, and numbers are the shapes of the tensors. C denotes the number of input features. ELU stands for the exponential linear unit (Clevert et al., 2015).

We split the 31 glaciers into six folds (Fig. 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*, *water* and *cloud*, which were assumed to be correctly labelled by the experts and are not glacier facies. 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 $s = \max_i \widehat{\mathbf{y}}_i$ of the predicted class and whether the prediction was correct ($\mathrm{acc} = \mathbb{1}\{\arg\max_i \widehat{\mathbf{y}}_i = \arg\max_i \mathbf{y}_i\} \in \{0,1\}$). 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 $s \mapsto \Pr(\mathrm{acc} = 1 \mid s)$, 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 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 $\beta_{0...3}$ 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 $\beta_{0...1,i}$ 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. Across the different feature sets, overall accuracy in terms of the macro-average F_1 score varied before confident learning, with no strict monotonic trend, and individual classes responded differently to the added predictors. 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 4). Under this evaluation, feature set 4 achieved the highest macro-average F_1 score at 74%, outperforming the other feature sets (71%, 72% and 72%,

Table 2. Out-of-sample test performance before confident learning. Bold font marks column bests. Aggregated metrics are shown in a larger font.

Feature set	Precision, %									Recall, %									F₁ score, %								
	ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	ice	snow	debris	firn	shadow	refrozen	water	cloud	Average
1	88	85	84	54	85	56	47	51	**69**	82	72	90	67	94	**79**	24	68	**72**	85	78	87	60	89	66	32	58	69
2	**90**	82	85	59	82	**61**	43	22	66	**85**	77	84	64	**96**	74	23	72	72	**87**	79	85	61	88	**67**	30	34	67
3	89	83	81	**59**	92	37	29	**72**	68	78	**80**	**94**	63	95	59	**68**	65	68	83	**82**	**87**	61	94	46	**41**	**68**	**70**
4	76	**86**	**92**	58	**94**	23	**50**	40	67	80	77	56	**71**	94	68	21	**74**	68	78	81	70	**64**	**94**	34	29	52	63

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*, *water* and *cloud* were not pruned, and thus the corresponding pruning rates are not reported.

Feature set	Precision, %									Recall, %									F₁ score, %									Pruned, %					
	ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	ice	snow	debris	firn	refrozen	Overall
1	95	94	90	78	66	65	52	86	78	88	93	76	87	96	84	30	79	79	92	93	83	82	78	73	38	82	78	11	20	3	22	13	13
2	94	96	94	85	69	84	62	77	83	96	92	79	86	96	88	23	75	79	95	94	86	85	81	86	33	76	80	10	20	2	24	18	12
3	95	97	85	79	85	88	69	32	79	87	88	94	88	96	85	42	79	82	91	92	89	82	90	87	53	45	79	13	17	3	23	29	13
4	95	93	87	78	95	76	65	82	82	89	92	96	84	95	81	47	67	81	92	92	91	81	95	78	55	68	82	14	17	27	18	24	16

respectively) and the models before confident learning. This corresponds to an improvement of about $0 \ldots 24\%$ in F₁ score for each class compared to the model without confident learning. In other words, incorporating the full suite of features together with confident learning yielded the best overall classification performance.

225 Using feature set 4 after confident learning as the best one in general, we analysed the confusion matrix (Fig. 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 classes—*ice* (F₁ score $= 92\%$), *snow* (92%), *debris* (91%), *firn* (81%) and *shadow* (95%)—were classified with high accuracy, as also indicated by high percentages along the matrix diagonal for these classes. However, the matrix also revealed systematic con-

230 fusion 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 observed between *firn* vs. *ice*, reflecting the gradual transition from older firn to bare glacier ice, as well as between *ice* and *refrozen-like* 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 other dark targets (*debris* or *shadow*) or as *refrozen-like*, as the latter closely resembles slush in satellite

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

Feature set	Precision, %									Recall, %									F1 score, %								
	ice	*snow*	*debris*	*firn*	*shadow*	*refrozen*	*water*	*cloud*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*refrozen*	*water*	*cloud*	Average	*ice*	*snow*	*debris*	*firn*	*shadow*	*refrozen*	*water*	*cloud*	Average
1	91	87	84	61	65	57	41	**79**	71	81	**83**	74	73	96	73	30	**79**	74	86	**85**	79	67	77	64	35	**79**	71
2	91	89	87	**66**	69	75	45	60	73	**89**	80	77	71	**96**	**73**	23	75	74	**90**	84	82	**70**	80	**74**	30	67	72
3	**92**	**90**	78	61	84	**77**	52	31	70	79	76	92	**77**	96	70	42	**79**	**76**	85	83	85	68	89	73	47	44	72
4	91	87	**88**	60	**93**	53	**59**	56	73	83	79	**92**	74	95	65	**47**	67	76	87	83	**90**	66	**94**	58	**52**	61	**74**

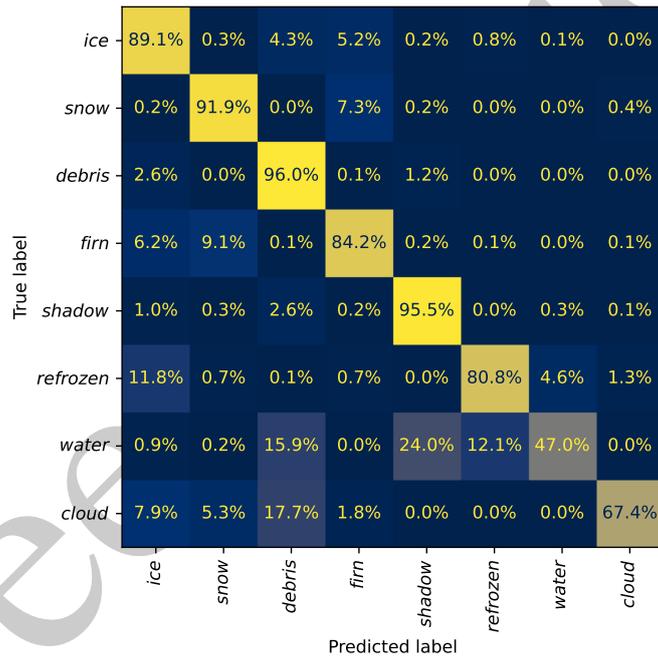**Figure 4.** Confusion matrix for the best classification results obtained after confident learning, normalised per row.

images. Another notable pattern was the misclassification of approximately one third of *cloud* pixels. The models attributed most of these pixels to the classes visible under the clouds. As a result, a portion of 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. To assess the robustness of our method across contrasting glacierised landscapes, we evaluated the models separately for five regions (Table 5). The models demonstrate a solid performance (F_1 score $\geq 74\%$) for all cases with reasonable sample size ($n > 500$), and less stable behaviour for the rest ($0\% \leq F_1$ score $\leq 100\%$). Notably, the sample size (n) shows moderate correlation with F_1 score ($r_s = 0.52, p \ll 0.01$). Similarly, the performance across different sensors remains consistent and correlated with the sample size (Table 6, $r_s = 0.41, p = 0.0169$). Glacier-wise macro-average F1 scores are $82.3 \pm 10.5\%$.

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

Region	Precision, % ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	Recall, % ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	F₁ score, % ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	n ice	snow	debris	firn	shadow	refrozen	water	cloud	Overall
Pyrenees	100	–	–	–	–	–	–	–	100	100	–	–	–	–	–	–	–	100	100	–	–	–	–	–	–	–	100	53	0	0	0	0	0	0	0	53
Alps	96	85	90	90	94	29	0	98	73	97	92	97	82	92	21	0	81	70	96	88	94	86	93	24	0	89	71	10448	4081	1634	5516	1414	19	4	101	23217
Scandinavia	92	74	58	91	97	83	–	85	83	97	89	78	76	98	62	–	94	85	95	81	66	83	97	71	–	89	83	3653	2283	49	3738	420	203	0	18	10364
Iceland	96	100	87	68	82	51	78	13	72	85	90	96	92	83	81	73	11	76	90	94	91	78	82	62	76	12	73	22177	11757	13734	5655	1046	404	252	103	55128
Svalbard	95	95	84	66	97	90	53	75	82	88	97	88	83	98	85	30	77	81	92	96	86	74	97	88	39	76	81	7883	6325	777	2204	7046	1166	391	400	26192

Table 6. Per-sensor performance obtained with the best fold models.

Platform	Precision, % ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	Recall, % ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	F₁ score, % ice	snow	debris	firn	shadow	refrozen	water	cloud	Average	n ice	snow	debris	firn	shadow	refrozen	water	cloud	Overall
Landsat 5	96	95	95	91	95	97	6	73	81	96	95	98	89	93	80	67	69	86	95	96	96	90	94	88	11	71	80	8443	4919	4033	2716	1542	282	3	264	22202
Landsat 7	94	92	98	71	83	67	95	89	86	68	87	91	94	94	80	69	94	85	79	90	94	81	88	73	80	92	85	2898	1393	375	2756	126	50	29	18	7645
Landsat 8	95	97	82	78	96	89	84	77	87	88	91	94	88	98	91	71	80	88	91	94	88	83	97	90	77	79	87	13600	6179	5540	4303	3372	722	207	250	34173
Landsat 9	100	–	–	–	–	–	–	–	100	100	–	–	–	–	–	–	–	100	100	–	–	–	–	–	–	–	100	15	0	0	0	0	0	0	0	15
Sentinel-2	95	90	86	76	94	60	57	26	73	90	92	97	76	94	71	33	22	72	93	91	91	76	94	65	42	24	72	19258	11955	6246	7338	4886	738	408	90	50919

We calibrated the model confidence scores to make them match real accuracies more closely. Before this step, the expected calibration error (ECE_{100}) was 11.6% (between $8\ldots18\%$ and between $10\ldots19\%$ across the test and validation folds, respec-
245 tively). After the calibration procedure, the error fell to 4.0% ($6\ldots17\%$ across test folds and $1\ldots5\%$ across validation folds). Figure 5 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 cali- brated predictive confidence. Calibration error, evaluated as the alignment between confidence and accuracy among samples
250 predicted as a given class, remains satisfactory for *ice* ($ECE_{100} = 5.6\%$), *snow* (6.3%), *debris* (6.2%), and *shadow* (6.7%), but exceeds 10% for *firn* (24%), *refrozen-like* (30%), *water* (18%) and *cloud* (20%).

Qualitative evaluation of the classification maps (Fig. 6) 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 ex-
255 ample, along the *snow-firn* and *firn-ice* interfaces—consistent with the inherently gradual transitions and hence higher ex- pert disagreement in those areas (Fig. 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 plausi- bly (Fig. 6a–j). Two notable failure modes were identified. In one case (the Langjokull 2000 scene, Fig. 6k), the model facies
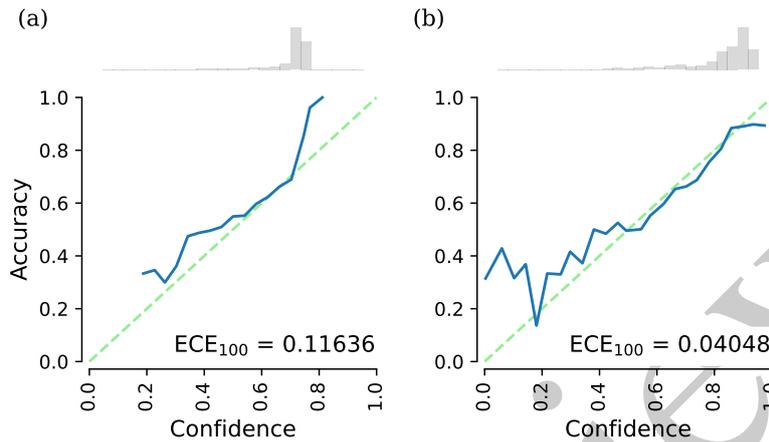
Figure 5. (a) Pre- and (b) post-calibration curves for the complete out-of-sample test data. The upper panels show the predictive confidence frequency. ECE_{100} stands for expected calibration error.

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, Fig. 6l), 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 (Fig. 8a, Equation 2), the facies data yielded a moderate, yet statistically significant Pearson correlation of $r = 0.65$ with measured SMB and an RMSE of 0.60 m w.e. When we calibrated a separate model for each glacier (Fig. 8b, Equation 3), the correlation improved to $r = 0.79$ and RMSE = 0.28 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, validation protocols and our limited sample size ($n = 9$). 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.
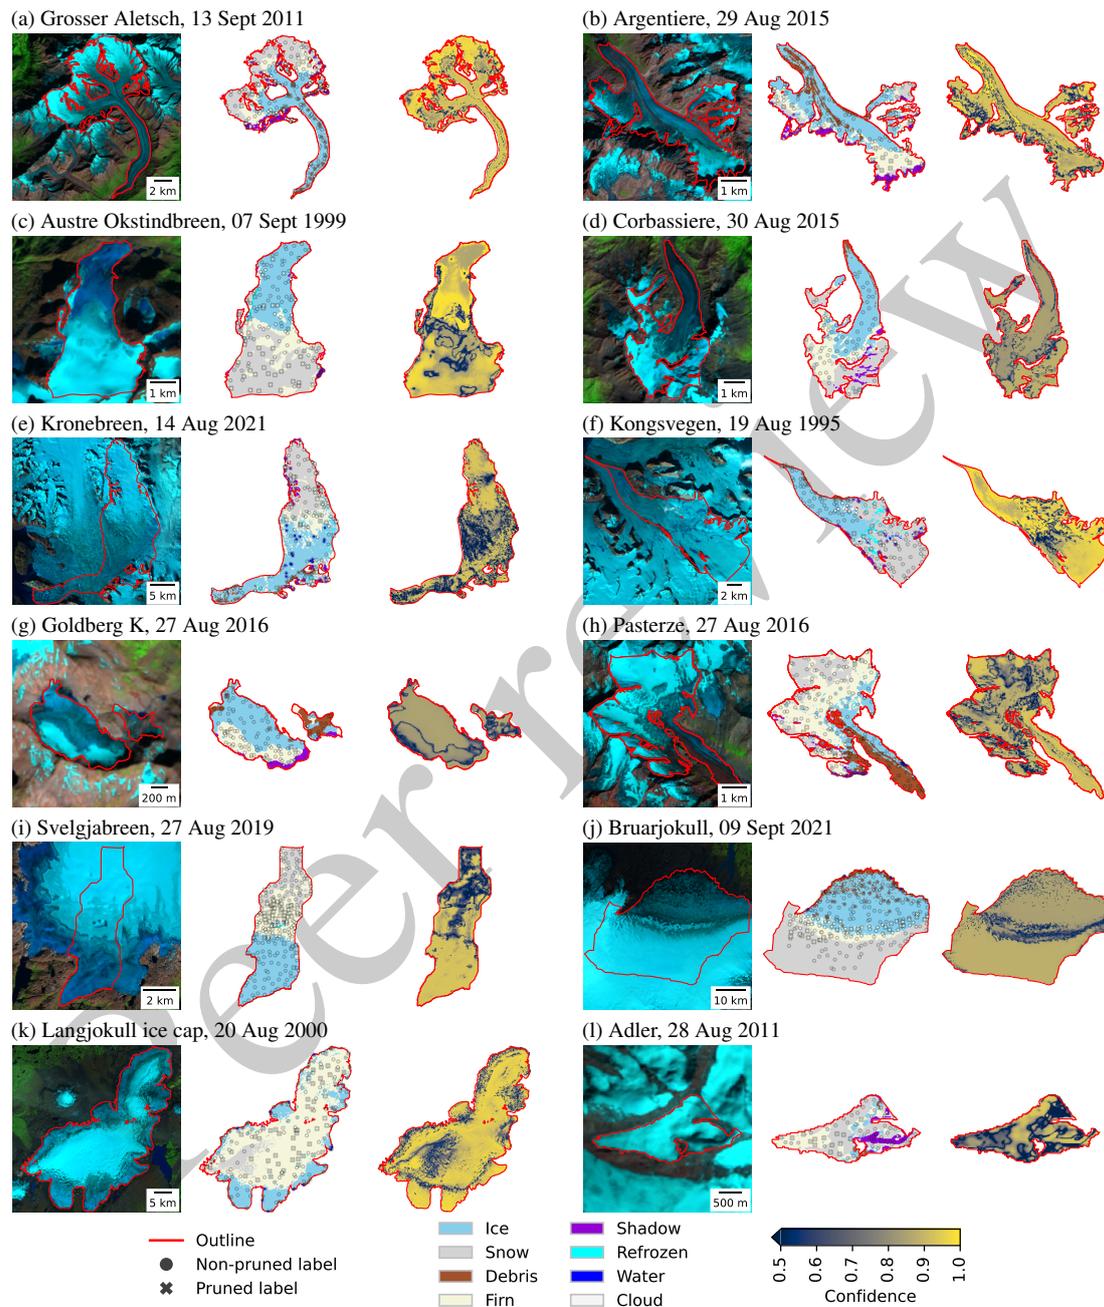
13

Figure 6. 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: $SWIR_{\approx 2.2\mu m}$, G: NIR, B: R). Landsat images courtesy of the U.S. Geological Survey. Copernicus Sentinel data 2015, 2016, 2019 and 2021.
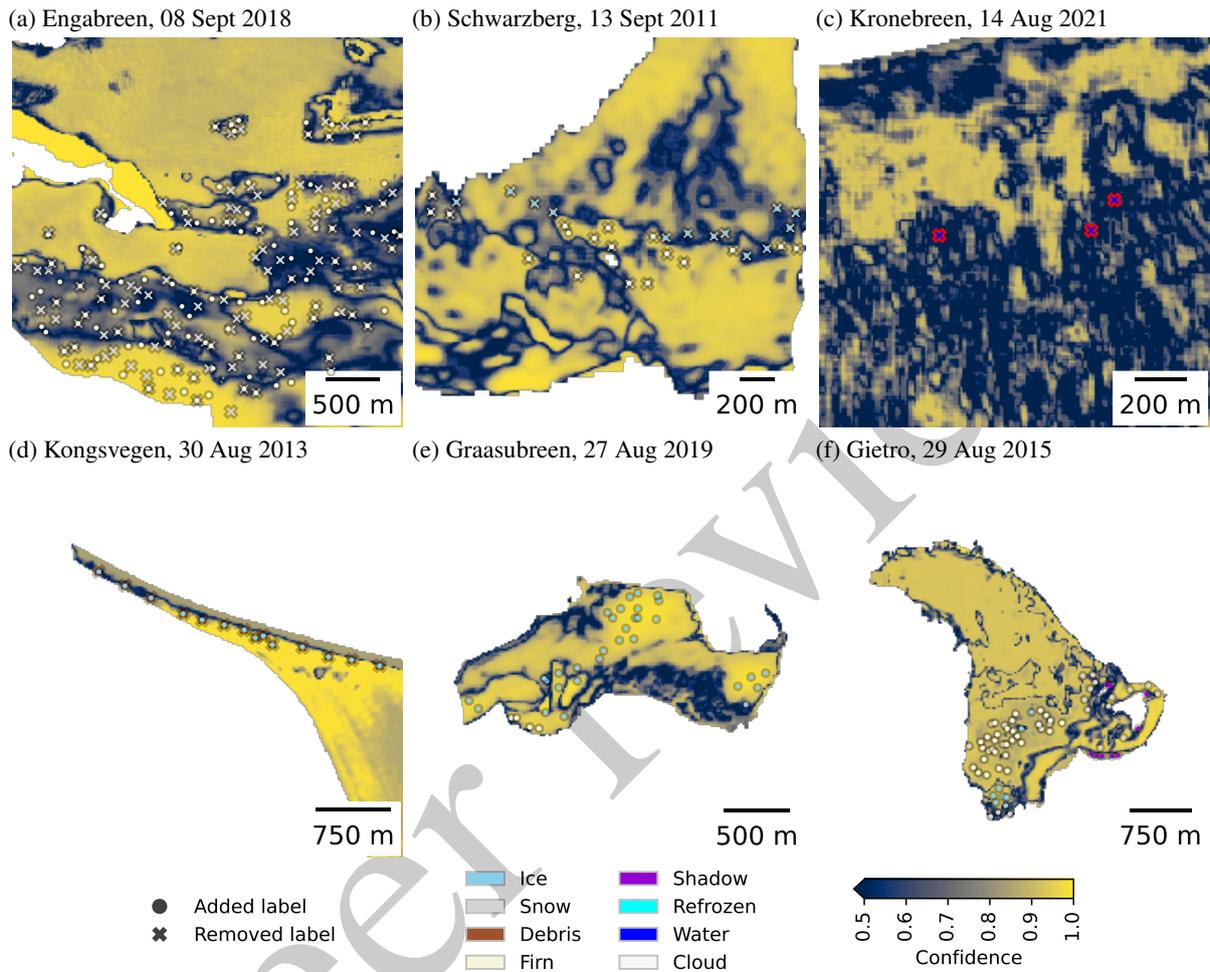
Figure 7. The same as in Fig. 2, but shown against predictive confidence.

5 Discussion

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 spatial extent, and we provide a baseline model that 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 (82%) 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

15
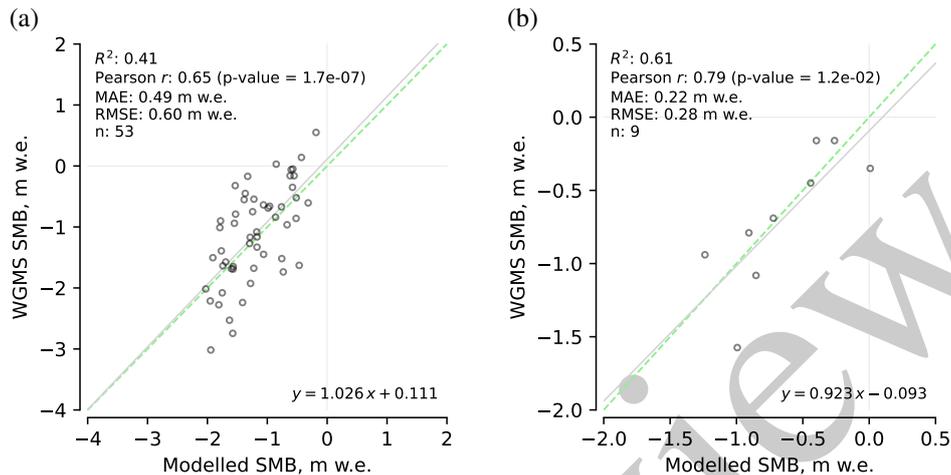
Figure 8. Glacier-wide SMB modelling results (a) of the six-fold cross-validation of the global model and (b) of the leave-one-year-out validation of the glacier-specific models for Kongsvegen and Midtre Lovenbreen.

285 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 a solid foundation for 290 future glacier mapping work.

The regional analysis confirms that model robustness depends on data availability. The positive correlation between sample size and F_1 score ($r_s = 0.52, p \ll 0.01$) 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 fewer examples occasionally show degraded results. Nonetheless, the glacier-wise macro-average F_1 score of $82.3 \pm 10.5\%$ 295 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.

300 Despite these positive results, several limitations are acknowledged. Glacier facies boundaries are inherently mixed 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

16

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 greater, 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 issue: 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. Further efforts should also focus on adding more labels for heavily underrepresented classes, namely *refrozen-like*, *water* and *cloud*. 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 0.28 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 0.60 m w.e. This error is admittedly 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 $\approx 0.604\ldots1.071$ 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 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 consid-

17

340 erable potential in integrating the CNN-derived facies products and their estimated uncertainty into physics-informed SMB frameworks, enhancing transferability across diverse climatic regimes.

6   Conclusions

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

345 via confident learning, we achieved accurate facies maps for diverse glaciers (macro-average F_1 score $= 82\%$ after pruning). To our knowledge, this is 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.

References

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

Ali, I., Shukla, A., and Romshoo, S. A.: Assessing linkages between spatial facies changes and dimensional variations of glaciers in the upper Indus Basin, western Himalaya, Geomorphology, 284, 115–129, https://doi.org/https://doi.org/10.1016/j.geomorph.2017.01.005, 2017.

365 Anilkumar, R., Bharti, R., Chutia, D., and Aggarwal, S. P.: Modelling point mass balance for the glaciers of the Central European Alps using machine learning techniques, The Cryosphere, 17, 2811–2828, https://doi.org/10.5194/tc-17-2811-2023, 2023.

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

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

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

Clevert, D. A., Unterthiner, T., and Hochreiter, S.: Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs), in: 4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings, https://arxiv.org/abs/1511.07289v5,

375 2015.

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

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

Drolon, V., Maisongrande, P., Berthier, E., Swinnen, E., and Huss, M.: Monitoring of seasonal glacier mass balance over the European Alps

380 using low-resolution optical satellite images, Journal of Glaciology, 62, 912–927, https://doi.org/10.1017/JOG.2016.78, 2016.

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

Florath, J., Keller, S., Staub, G., and Weinmann, M.: 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:

385 Evolution in Remote Sensing (WHISPERS), pp. 1–5, https://doi.org/10.1109/WHISPERS52202.2021.9484055, 2021.

Fox, E., Schwartz-Marin, E., Rangecroft, S., Palmer, S., and Harrison, S.: Water resource framing for the value and governance of glacier water availability in the semi-arid Chilean Andes, Frontiers in Water, Volume 6, https://doi.org/10.3389/frwa.2024.1367889, 2024.

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

390 Gupta, R., Haritashya, U., and Singh, P.: Mapping dry/wet snow cover in the Indian Himalayas using IRS multispectral imagery, Remote Sensing of Environment, 97, 458–469, https://doi.org/https://doi.org/10.1016/j.rse.2005.05.010, 2005.

Hall, D., Ormsby, J., Bindschadler, R., and Siddalingaiah, H.: Characterization of Snow and Ice Reflectance Zones On Glaciers Using Landsat Thematic Mapper Data, Annals of Glaciology, 9, 104–108, https://doi.org/10.3189/S0260305500000471, 1987.

Hall, D. K., Chang, A. T., and Siddalingaiah, H.: Reflectances of glaciers as calculated using Landsat-5 Thematic Mapper data, Remote

395 Sensing of Environment, 25, 311–321, https://doi.org/https://doi.org/10.1016/0034-4257(88)90107-1, 1988.

Herreid, S. and Pellicciotti, F.: The state of rock debris covering Earths glaciers, Nature Geoscience, 13, 621–627, https://doi.org/10.1038/s41561-020-0615-0, 2020.

Hock, R., Kootstra, D.-S., and Reijmer, C.: 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, 318, pp. 163–170, https://www.scopus.com/inward/record.uri?eid=2-s2.0-38549156622&partnerID=40&md5=abe3ce04fbedff2850dd16badab023de, 2007.

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

Huo, D., Bishop, M. P., and Bush, A. B. G.: Understanding Complex Debris-Covered Glaciers: Concepts, Issues, and Research Directions, Frontiers in Earth Science, Volume 9, https://doi.org/10.3389/feart.2021.652279, 2021.

Jawak, S. D., Wankhede, S. F., Luis, A. J., and Balakrishna, K.: 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, https://doi.org/10.3390/rs14246311, 2022.

Kingma, D. P. and Ba, J. L.: Adam: A Method for Stochastic Optimization, in: 3rd International Conference on Learning Representations, ICLR 2015—Conference Track Proceedings, International Conference on Learning Representations, ICLR, https://doi.org/10.48550/arxiv.1412.6980, 2014.

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

König, M., Winther, J.-G., Kohler, J., and König, F.: Two methods for firn-area and mass-balance monitoring of Svalbard glaciers with SAR satellite images, Journal of Glaciology, 50, 116128, https://doi.org/10.3189/172756504781830286, 2004.

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

Liang, S.: Narrowband to broadband conversions of land surface albedo I: Algorithms, Remote Sensing of Environment, 76, 213–238, https://doi.org/https://doi.org/10.1016/S0034-4257(00)00205-4, 2001.

Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P.: Focal Loss for Dense Object Detection, in: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 2999–3007, https://doi.org/10.1109/ICCV.2017.324, 2017.

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

Loshchilov, I. and Hutter, F.: 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, https://arxiv.org/abs/1608.03983v5, 2016.

Lu, Y., Zhang, Z., Shangguan, D., and Yang, J.: Novel Machine Learning Method Integrating Ensemble Learning and Deep Learning for Mapping Debris-Covered Glaciers, Remote Sensing, 13, 2595, https://doi.org/10.3390/RS13132595, 2021.

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

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

Müller, F.: Zonation in the Accumulation Area of the Glaciers of Axel Heiberg Island, N.W.T., Canada, Journal of Glaciology, 4, 302–311, https://doi.org/10.3189/S0022143000027623, 1962.

Northcutt, C. G., Jiang, L., and Chuang, I. L.: Confident Learning: Estimating Uncertainty in Dataset Labels, Journal of Artificial Intelligence Research, 70, 1373–1411, https://doi.org/10.1613/JAIR.1.12125, 2021.

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

Podsiadlo, I., Paris, C., Callegari, M., Marin, C., Günther, D., Strasser, U., Notarnicola, C., and Bruzzone, L.: 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, https://doi.org/10.1109/JSTARS.2020.3028653, 2020.

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

Prieur, C., Rabatel, A., Thomas, J.-B., Farup, I., and Chanussot, J.: Machine Learning Approaches to Automatically Detect Glacier Snow Lines on Multi-Spectral Satellite Images, Remote Sensing, 14, https://doi.org/10.3390/rs14163868, 2022.

Rabatel, A., Dedieu, J. P., and Vincent, C.: Using remote-sensing data to determine equilibrium-line altitude and mass-balance time series: validation on three French glaciers, 19942002, Journal of Glaciology, 51, 539–546, https://doi.org/10.3189/172756505781829106, 2005.

Rabatel, A., Sirguey, P., Drolon, V., Maisongrande, P., Arnaud, Y., Berthier, E., Davaze, L., Dedieu, J. P., and Dumont, M.: 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, https://doi.org/10.3390/rs9050507, 2017.

Rastner, P., Prinz, R., Notarnicola, C., Nicholson, L., Sailer, R., Schwaizer, G., and Paul, F.: On the Automated Mapping of Snow Cover on Glaciers and Calculation of Snow Line Altitudes from Multi-Temporal Landsat Data, Remote Sensing, 11, https://doi.org/10.3390/rs11121410, 2019.

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

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

Simonyan, K. and Zisserman, A.: 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, https://doi.org/10.48550/arXiv.1409.1556, 2014.

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

Wendleder, A., Schmitt, A., Erbertseder, T., DAngelo, P., Mayer, C., and Braun, M. H.: Seasonal Evolution of Supraglacial Lakes on Baltoro Glacier From 2016 to 2020, Frontiers in Earth Science, Volume 9 - 2021, https://doi.org/10.3389/feart.2021.725394, 2021.

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

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

Winther, J.-G.: Landsat TM derived and in situ summer reflectance of glaciers in Svalbard, Polar Research, 12, 37–55, https://doi.org/https://doi.org/10.1111/j.1751-8369.1993.tb00421.x, 1993.

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

Yousuf, B., Shukla, A., Arora, M. K., and Jasrotia, A. S.: 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, https://doi.org/10.1177/0309133319840770, 2019.

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

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.,
485   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., and Team, T.: Community estimate of global glacier mass changes from 2000 to 2023, Nature, 639, 382–388, https://doi.org/10.1038/s41586-024-08545-z, 2025.

Zhang, J., Jia, L., Menenti, M., and Hu, G.: Glacier Facies Mapping Using a Machine-Learning Algorithm: The Parlung Zangbo Basin Case
490   Study, Remote Sensing, 11, https://doi.org/10.3390/rs11040452, 2019.