



Out-of-the-box calving front detection method using deep learning

Oskar Herrmann¹, Nora Gourmelon², Thorsten Seehaus¹, Andreas Maier², Johannes J. Fürst¹, Matthias H. Braun¹, and Vincent Christlein²

¹Institute of Geography, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

²Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

Correspondence: Oskar Herrmann (oskar.herrmann@fau.de)

Abstract. Glaciers across the globe react to the changing climate. Monitoring the transformation of glaciers is essential for projecting their contribution to global mean sea level (GMSL) rise. The delineation of glacier-calving fronts is an important part of the satellite-based monitoring process. This work presents a calving front extraction method based on the deep learning framework nnU-Net, which stands for no new U-Net. The framework automates the training of a popular neural network, called U-Net, designed for segmentation tasks. Our presented method marks the calving front in Synthetic Aperture Radar (SAR) images of glaciers. The images are taken by six different sensor systems. A benchmark dataset for calving front extraction is used for training and evaluation. The dataset contains two labels for each image. One label denotes a classic image segmentation into different zones (glacier, ocean, rock, and no information available). The other label marks the edge between the glacier and the ocean, i. e., the calving front. In this work, the nnU-Net is modified to predict both labels simultaneously. In the field of machine learning, the prediction of multiple labels is referred to as Multi-Task-Learning (MTL). The resulting predictions of both labels benefit from simultaneous optimization. For further testing of the capabilities of MTL, two different network architectures are compared, and an additional task, the segmentation of the glacier outline, is added to the training. In the end, we show that fusing the label of the calving front and the zone label is the most efficient way to optimize both tasks with no significant accuracy reduction compared to the MTL neural network architectures. The automatic detection of the calving front with a nnU-Net trained on fused labels improves from the baseline Mean Distance Error (MDE) of 753 ± 76 m to 541 ± 84 m. The scripts for our experiments are published on Gitlab (https://gitlab.cs.fau.de/ho11laqe/nnunet_glacier.git).

1 Introduction

Unlike the large majority of land-terminating glaciers, marine and lake-terminating (MALT) glaciers reach a water body at a low elevation. The contact surface is often referred to as the calving front. Their ice is lost by sub-marine melting or calving, i. e., ice that breaks off, gets de-connected, and starts to float freely in the form of icebergs or ice floes. Both processes, sub-marine melting and calving, determine the total frontal ablation, i. e., the mass loss at the calving front. Frontal ablation is often a dominant factor in the total mass budget of MALT glaciers (McNabb et al., 2015; Shepherd et al., 2018; Minowa et al., 2021). Besides its importance in the total glacier mass balance, the representation of processes controlling frontal ablation is currently a pressing task for numerical glacier models (Beer et al., 2021). Neglecting frontal ablation can introduce an important bias. Recinos et al. (2019) analyzed the impact on ice thickness reconstruction (based on mass conservation) in Alaska and



reported an underestimation of 19 % on regional and up to 30 % on glacier scales. Various successful approaches exist to parameterize frontal ablation for individual glaciers. Still, the implementation in large-scale or global models is limited by the amount and quality of measurements to constrain the models (Recinos et al., 2019). Thus, large-scale measurements (ideally time series) of frontal ablation are demanded by the modelling community (Recinos et al., 2021). Driving forces of frontal ablation are, on the one hand, ice flux (higher flux, e. g., due to bed lubrication by meltwater can trigger calving events) and, on the other hand, marine factors such as ocean temperature, fjord bathymetry, and sea ice/ice mélange conditions (Carr et al., 2014; Straneo et al., 2013). For example, the persistence of the ice mélange in front of the glacier can stabilize the calving front and affect the glacier dynamics. In contrast, the breakup of the ice mélange can lead to increased calving and ice flow at the glacier terminus (Amundson et al., 2010; Kneib-Walter et al., 2021; Rott et al., 2020). Moreover, a significant frontal retreat can also indicate a retreat of the grounding zone (Friedl et al., 2018). The retreat of the grounding zone of a glacier with retrograde bedrock formation will lead to further grounding zone retreat, resulting in increased ice loss and destabilization of the glacier or ice stream (Robel et al., 2016). Thus, information on the temporal variability of the calving front position provides fundamental information on the state of the glacier or ice stream. Therefore, the glacier area has been defined as an Essential Climate Variable (ECV) product by the World Meteorological Organization (WMO). Calving front positions were usually manually mapped using different remote sensing imagery (Baumhoer et al., 2018). Only a few studies applied automatic or semi-automatic approaches. In polar regions, the ocean downstream of the glaciers is often covered by sea ice and calved off icebergs, forming the so-called ice mélange, making calving front delineation a challenging task, even when captured by hand. Deep learning approaches have shown high potential for carrying out such complex segmentation tasks, e. g., on medical imagery (Jang and Cho, 2019). In recent years, the application of deep learning techniques for glacier front detection started (Zhang et al., 2019; Cheng et al., 2021; Baumhoer et al., 2019; Hartmann et al., 2021). Most of these methods use optical imagery, which has the drawback of being dependent on daytime and cloud coverage. Additionally, most of these published methods are incomparable since the algorithms or the input data remained undisclosed.

The dataset published by Gourmelon et al. (2022) provides 681 SAR images of calving fronts. SAR imagery is independent of sunlight and cloud coverage, enabling continuous temporal coverage of the observation area, but compared to optical data, it has only one channel and has more speckle noise. For every SAR image, two labels are provided. One label provides four classes: ocean, glacier, rock, and no information available (e. g., radar shadow and layover areas, areas outside the swath). The other label marks the calving front with a one-pixel wide line. Based on the training set of the dataset, Gourmelon et al. (2022) train a modified U-Net for each label. One U-Net solves the task of glacier segmentation, and one detects the calving front. On the two test glaciers of the dataset, the segmentation model achieves an Intersection over Union (IoU) of 67.7 ± 0.6 . By taking the boundary between ocean and glacier, they extract a prediction of the calving front from the zone prediction with a Mean Distance Error (MDE) of 753 ± 76 m. The model trained directly on the front label achieved an MDE of 887 ± 189 m.

In this work, we present a method that utilises both labels of the dataset instead of training separate models for each task (see Fig. 1). Multi-Task-Learning (MTL) is a technique for machine learning algorithms that uses one model to tackle multiple tasks. In most cases, the potentially larger dataset and higher information content lead to higher performance for the individual tasks (Bischke et al., 2019; He et al., 2021; Li et al., 2019; Amyar et al., 2020; Chen et al., 2019).




Figure 1. Illustration of the modified nnU-Net for simultaneous predicting landscape zones and the glacier front. On the left is an exemplary satellite image of the Crane Glacier taken by TanDEM-X (TDX) on 24 June 2011. On the right, the upper image shows the four segmentation classes: ocean (white), glacier (light grey), rock (dark grey), and no information available (black). The lower image shows the glacier's calving front as a white line versus the black background.

Our method is based on the nnU-Net (no new U-Net) proposed by Isensee et al. (2021), which is an out-of-the-box framework for training the U-Net. The framework takes a fingerprint of the dataset and adjusts the hyperparameters accordingly. Therefore, the nnU-Net is a powerful tool that simplifies the application of deep learning algorithms. Its performance on segmentation of SAR data has not been tested, and the availability of two labels suggests the modification of the nnU-Net for MTL. As a
65 baseline for our evaluation of the nnU-Net, we use the results of Gourmelon et al. (2022).

In particular, our contributions are as follows: (1) Visual preparation of the dataset by Gourmelon et al. (2022). (2) Out-of-the-box application and evaluation of the nnU-Net for calving front detection and zone segmentation. (3) Evaluation of two different U-Net architectures for MTL. (4) Test if an artificial third label improves the calving front detection. (5) Introduction of an efficient MTL approach that fuses the two labels of the dataset. (6) Analysis of the influence of season, glacier, and
70 satellite on the performance of the model.

The paper is organized as follows: After presenting the related work in Sec. 2, we give an overview of the dataset (Gourmelon et al., 2022) in Sec. 3. Section 4 explains the method and our six experimental setups. Section 5 examines the results and analyses the influence of different properties of the satellite images. Section 6 summarises the work.

2 Related Work

75 Automated monitoring of glacier-covered areas is a growing research field. Recent glacier monitoring uses deep learning methods due to the increasing availability of satellite images and computing power. Many methods are based on the U-Net (Ronneberger et al., 2015). One approach is to segment the images into different areas and extract the calving front as the



border between segmentation areas (Hartmann et al., 2021; Zhang et al., 2019; Baumhoer et al., 2019; Periyasamy et al., 2022; Loebel et al., 2022). Another approach directly trains a model on the position of the calving front (Davari et al., 2022). This task suffers from severe class imbalance due to the thin calving front. They approach this problem by creating a distance map from every pixel to the front line. The network is trained on the distance map instead of the thin front line. The actual front-line prediction is then extracted during post-processing.

Other works use the segmentation network DeepLabv3 (Chen et al., 2018) to detect calving fronts. The main advantage of DeepLabv3 over U-Net is the atrous spatial pyramid pooling, which makes the network adaptable to different image resolutions (Zhang et al., 2021; Cheng et al., 2021).

Calving Front Machine (CALFIN) proposed by Cheng et al. (2021) segments optical and SAR images into the ocean-land zones and extract the calving front during post-processing with a topography map of the area. They apply MTL with a late-branching architecture. They use two labels: a binary ocean mask and a binary calving front mask. They achieve state-of-the-art predictions with an 86 m deviation from the measured calving front. A detailed comparison to aforementioned U-Net-based methods by Mohajerani et al. (2019) and Baumhoer et al. (2019) revealed the generalization ability of CALFIN on other glacier datasets. The images had to be down-sampled for CALFIN. Therefore, the distance in meters doubles but comparing the pixel distance errors reveals a similar performance. With 29 M parameters, CALFIN is still a large network that needs a large amount of training data. Chen et al. (2019) propose a similar approach for medical image segmentation and show that the individual tasks benefit from MTL. Several other approaches have advanced the U-Net architecture for MTL for medical applications (Abolvardi et al., 2020; Kholiavchenko et al., 2020; Li et al., 2019; Amyar et al., 2020). The work of Heidler et al. (2021) uses MTL for the segmentation of a binary ocean-land mask and edge detection of the Antarctic coastline, where the calving front is just part of the coastline. They add task-specific heads for the two tasks to the U-Net. They achieve results with a deviation of coastline of 345 m compared to 483 m with the vanilla U-Net. Their post-processing involves the terrain elevation, and the metric is calculated within 2 km of the true coastline.

100 3 Dataset

The dataset used in this work is provided by Gourmelon et al. (2022). It contains 681 Synthetic Aperture Radar (SAR) images of seven marine-terminating glaciers taken by six different satellites. Two glaciers are located in the northern hemisphere, namely Columbia Glacier in Alaska and Jakobshavn Isbrae (Sermeq Kujalleq) in Greenland. The five glaciers in the southern hemisphere are all located on the Antarctic Peninsula (AP) (see Fig. 2). The Crane, Mapple, and Jorum glaciers are closest to the south pole, followed by Dinsmoore, Bombardier, Edgeworth. They are located so close together that they are treated as one glacier site. The last glacier is the Sjögren-Inlet (SI) Glacier.

Table 1 lists the seven glacier sites with the number of images and the area they show. The table also depicts the train-test split. The samples of the train set are used for the parameter optimization of the segmentation method, and the test set is exclusively used for evaluating the method. The test set contains the glaciers Mapple and Columbia. Columbia is the only mountain glacier in the dataset. The other glaciers are outlet glaciers of polar ice sheets. Therefore, Columbia can be seen as a

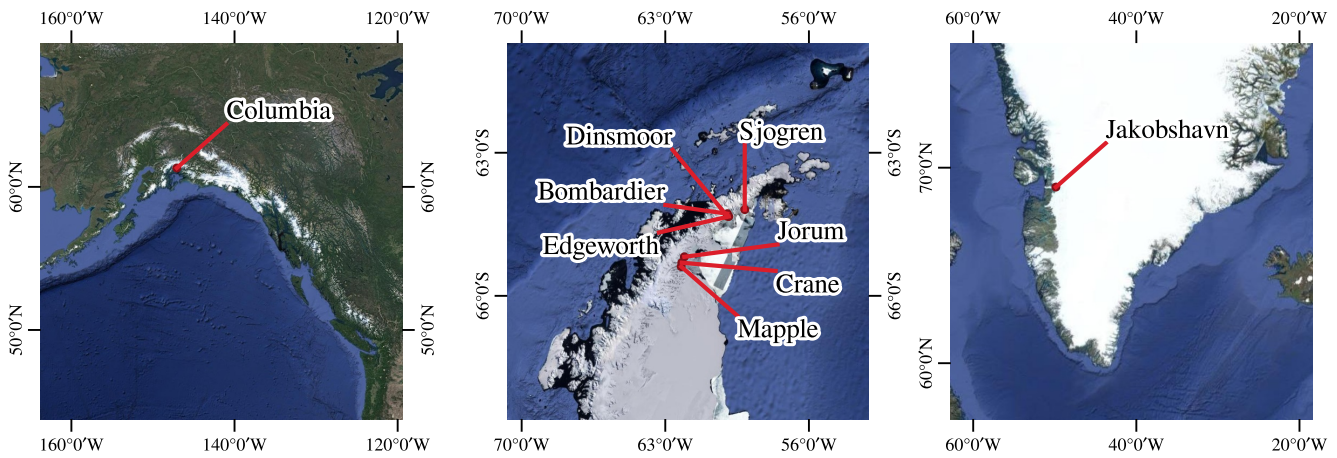


Figure 2. Location of the seven benchmark glaciers: Columbia Glacier in Alaska; Crane; Mapple; Jorum; Dinsmoore, Bombardier, Edgeworth (DBE); Sjøgren Glacier in the Antarctic Peninsula; and Jakobshavn Isbrae Glacier in Greenland. The background images are taken from the Google Maps Satellite imagery layer © Google 2019.

	Alaska	Antarctic Peninsula					Greenland
	Columbia	Mapple	Crane	Jorum	DBE	Sjøgren-Inlet	Jakobshavn Isbrae
Split	- test: 122 -		- train: 559 -				
Images #	65	57	69	77	133	121	159
Area [km]	32 x 15	8 x 8	19 x 25	20 x 13	22 x 20	23 x 19	16 x 19

Table 1. Properties of the dataset, including a list of captured glaciers, train-test split, number of images per glacier, and covered area.

benchmark glacier for the generalization capability of the segmentation method. Columbia and Mapple also differ in the shape of the glacier. The images of Columbia show multiple calving fronts in one image, and the flow of the glacier arms goes in different directions. On the other hand, Mapple is a less complex glacier, moving in only one direction with one clearly defined calving front between the lateral fjord side walls.

115 The dataset contains two labels and a bounding box of the glacier for each SAR image. One shows a mask of the different zones of the glacier (ocean, glacier, no information available, rock). The other label contains a one-pixel wide line representing the calving front. The bounding box is not utilized for our method. A sample of each glacier in the training set with its corresponding labels is shown in Fig. 3. Predicting the zone mask can be seen as a classic segmentation problem. The calving front can also be extracted from the zone label by taking the border between ocean and ice areas in the corresponding bounding
 120 box. The direct delineation of the calving front is a more difficult task due to the high-class imbalance. Fewer than 1 % of the pixels are labelled as front pixels. Additionally, the structure of the class region is not a convex hull but a thin line.

Corresponding to the change of glacier area, the position of the calving front changes. In Fig. 4, the retreat of Columbia is visualized by plotting the calving front position of every sample in a different colour. The bright lines represent past calv-

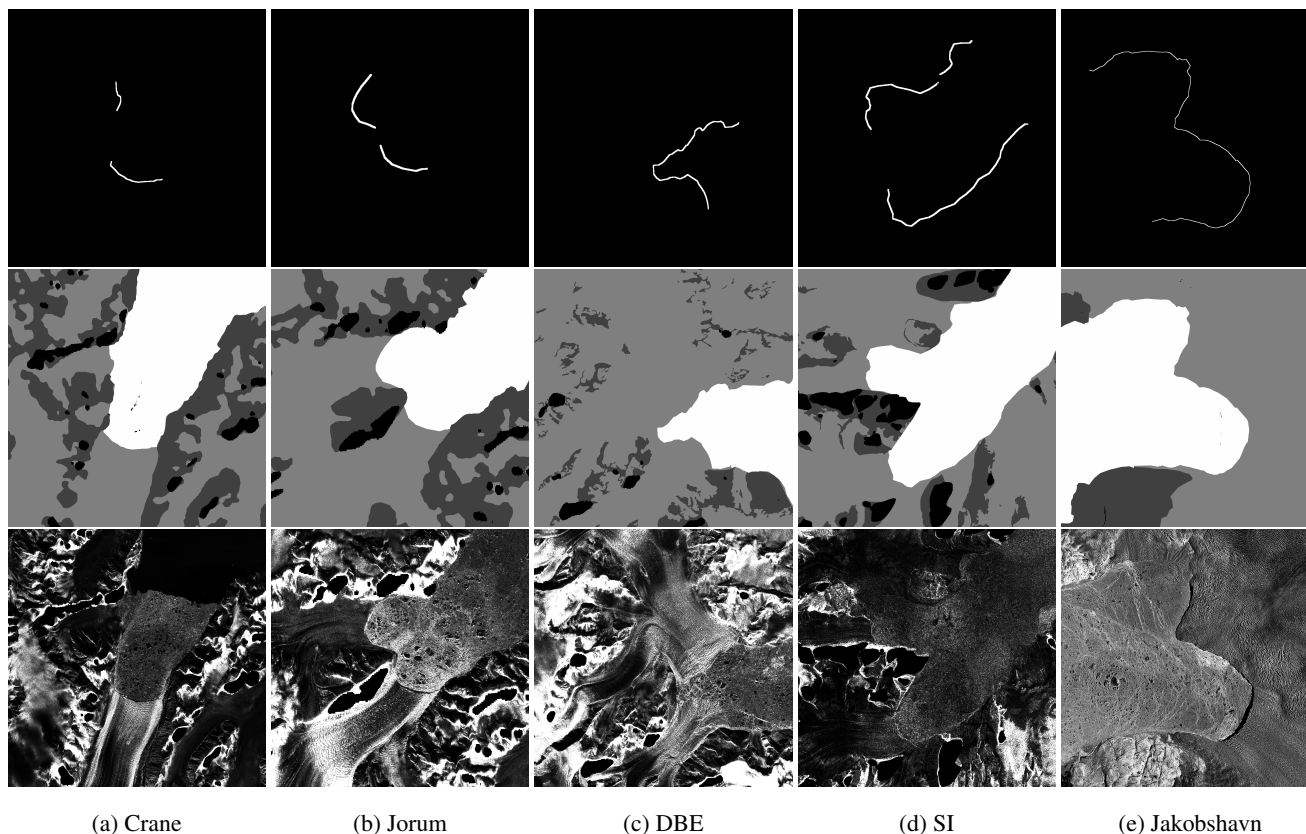


Figure 3. Sample images of every glacier in the training set and their corresponding labels. The first row shows the front label with a black background and a one-pixel wide white line representing the calving front. The second row contains the zone labels with four classes: ocean (white), glacier (light grey), rock (dark grey), and no information available (black). SAR imagery is provided by DLR, ESA, and ASF.

ing fronts, and the dark red lines the more recent positions. The constant retreat of Columbia (COL) is visible through the
125 tiered lines. The visualization of the glaciers in the southern hemisphere is included in Appendix B. The change in the class
distribution of the zone labels over time is shown in Appendix. C.

Every glacier is captured by multiple satellites for a higher temporal resolution and extended observation periods. Meaning
that recordings of one glacier are captured by different SAR systems with different image resolutions. In Fig. 5, a timeline of
the images of each glacier visualizes the observation time and frequency of the images. The first two rows show the glaciers of
130 the test set.




Figure 4. Evolution of the calving front position of Columbia Glacier (Alaska). The red lines are all calving front labels available in the dataset. Background: SAR intensity image acquired by TDX on date 13 November 2011. Projection CRS EPSG:4326 - WGS 84 / UTM zone 6N. SAR imagery is provided by DLR, ESA, and ASF.



Figure 5. Distribution of the dataset's images over time. The samples are grouped by the seven glaciers and coloured according to the capturing satellite.



4 Method

In this section, we explain our method, which includes the utilisation of the nnU-Net as a framework that simplifies the training of a U-Net. We document our six experimental setups that aim to evaluate the impact of MTL on the training of the U-Net.

4.1 Background

135 In the field of deep learning, a lot of time and effort is put into the hyperparameter search. Isensee et al. (2021) proposes the nnU-Net that automates the manual tuning of hyperparameters. It takes a fingerprint of the dataset and adjusts hyperparameters accordingly. The nnU-Net is designed with a focus on three-dimensional (3D) data like Computer Tomography (CT) and Magnetic Resonance Imaging (MRI), which have different spacing in dimensions. One parameter of the data fingerprint is the distribution of spacing. The annotation, image resampling strategy, and target spacing depend on the spacing distribution
140 as rule-based parameters. Two other fingerprint parameters are intensity distribution and image modality, which determine intensity normalization. If the imaging modality is CT, a global dataset percentile clipping and z-score normalization are applied. For most samples, a large portion of the image is unlabeled, which is called the background. All the labelled areas are called foreground. For the normalization, the background is omitted, and the mean and standard derivation is calculated for the foreground. In all other cases, the z-score normalization is applied with per image mean and standard deviation. The
145 last fingerprint parameter is the median shape of the samples. The shape and the available Graphics Processing Unit (GPU) memory determine the patch and batch size and, therefore, the network topology. The patch size is initialized with the median shape and iteratively reduced while adapting the network topology accordingly until the network can be trained with a batch size of at least two given GPU memory constraints.

In addition to the rule-based parameters, there are fixed parameters. These parameters are based on the authors' experience
150 and generalize well across various tasks. The nnU-Net uses a poly learning rate scheduler, a combination of dice coefficient and cross-entropy as the loss function, and Stochastic Gradient Descent (SGD) with a Nesterov momentum as optimizer. Additionally, deepsupervision is used to avoid vanishing gradients in neural networks with many layers. Independent of the dataset size, one epoch is defined as 250 mini-batches with foreground oversampling.

This framework achieves robust results for various medical image segmentation tasks. Overall, nnU-Net sets a new state of
155 the art in 33 of 53 segmentation tasks of the Kidney and Kidney Tumor Segmentation Challenge challenge (Heller et al., 2021) and otherwise shows performances on par with or close to the top leaderboard entries. But the performance of the nnU-Net on segmenting glacier SAR images is yet to be tested.

4.2 Multi-Task Learning with nnU-Net

The first research goal is to apply the nnU-Net out-of-the-box on the glacier front detection and on the glacier zone segmen-
160 tation, respectively. Training the nnU-Net directly on the front labels is the most straightforward approach for calving front detection. The nnU-Net is intended to be used in the STL manner. In Fig. 6, these two baseline experiments are represented by the two blue columns on the left. The label of the calving front is dilated to the width of five pixels. Our preliminary ex-

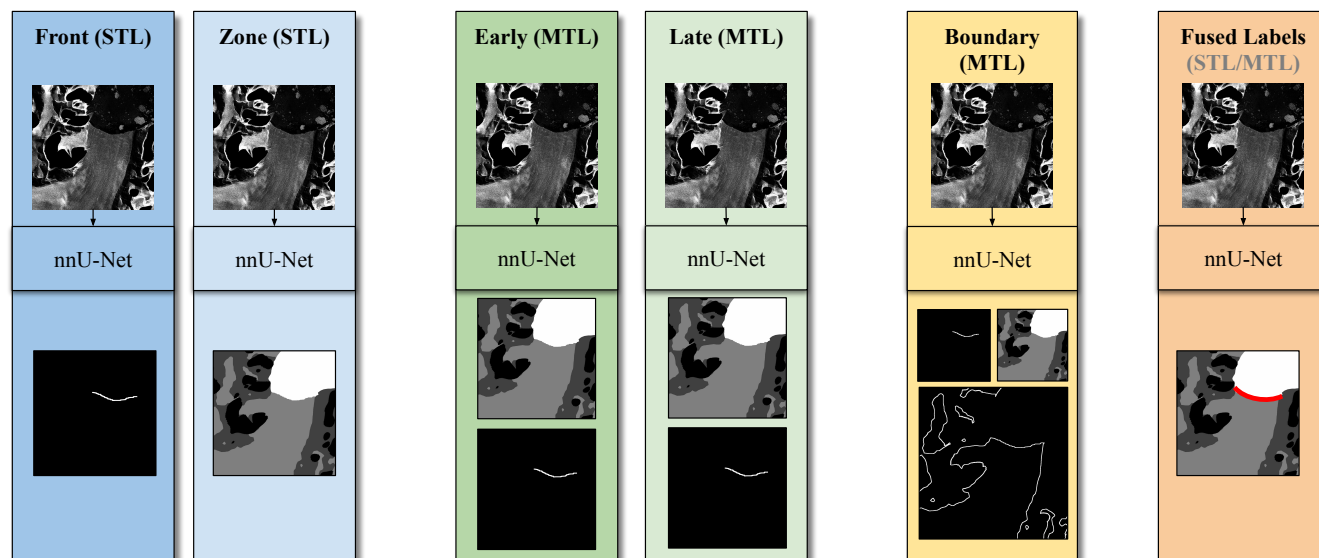


Figure 6. Illustration of the six experiments: Two Single-Task-Learning (STL) experiments (blue), two MTL experiments with different architectures (green), one MTL experiment with an additional task of delineating the whole glacier boundary (yellow), and one experiment where the two labels are fused into one segmentation task with an additional class in the zone label (orange).

periments have shown that dilation makes the predictions more robust. For the training with zone labels, the post-processing includes extracting the boundary between the ocean and the glacier.

165 In the following experiments, the segmentation problem changes to a multi-task problem, where both labels are used to train one model. The next two experiments concern network architecture. They are represented by the two green columns in Fig. 6. The early-branching architecture uses one decoder for every label. Thus, the number of parameters increases by about 50%. In contrast, the late-branching architecture requires only a small change to the vanilla U-Net. An additional channel of the last layer is used to predict the second label. For this architecture, only the weights of one kernel are trained additionally to the set
170 of parameters needed for one task. The change in the total number of parameters that need to be trained is negligible.

Because architecture changes with multi-task learning were not foreseen in nnU-Net framework, we had to make changes to the framework. During the experiment's planning, i. e., pre-processing phase, we fixed the network architecture's estimated size so that late- and early-branching networks return the same value for the network size. Otherwise, early and late branching networks would be trained on different patch sizes. Thus, performance differences during the evaluation might arise from the
175 different patch sizes and not the differing architectures. During training, the error of all labels is calculated and summed up with equal weighting. Only minor changes had to be made to the inference part. After the test samples are divided into patches and fed through the network and the patch predictions are combined into the prediction of the whole image. The predictions of the zones are post-processed to get an additional result for the position of the calving front. All glacier pixels with a neighbouring pixel classified as the ocean are classified as glacier front to extract the glacier front from the zone predictions.

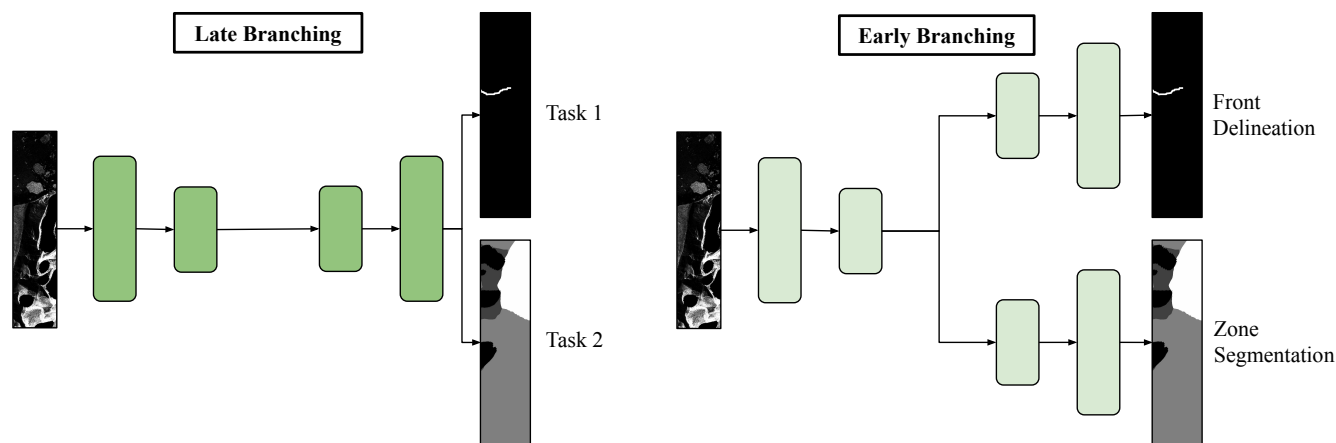


Figure 7. Illustration of early- and late-branching U-Net architectures for MTL. Both architectures perform joint feature extraction. The early-branching architecture applies separate reconstruction with two decoders, one for each task. The late-branching network applies joint reconstruction with a common decoder.

180 The last two experiments concern label changes. In Fig. 6, they are coloured yellow and orange. The fifth experiment of this work, cf. Fig. 6 (yellow), extracts the boundaries between the glacier zone and all other zones as a third segmentation task for the late-branching U-Net. The label of the glacier boundaries was extracted from the zone label. All glacier pixels with a neighbouring rock or shadow pixel are assigned as glacier boundaries. The hypothesis is that providing more information about the same sample benefits the performance of the U-Net on the individual tasks. The last experiment fuses the zone and
185 front labels by creating a fourth class in the zone label associated with the glacier front, cf. Fig. 6 (orange). As the front line has a width of five pixels (35 – 100 m depending on the image resolution), the other zone classes are merely impaired. A post-processing step is added to the predictions. The front pixels are isolated to generate a front prediction for comparing the results with the other experiments. To generate a prediction of the zones, pixels classified as front are assigned to the ocean, and the glacier zone is dilated once with a 7×7 kernel.

190 5 Evaluation

In this section, we will examine our evaluation metrics, compare the results of the six proposed experiments and evaluate the results of the fused label experiment. We used a five-fold cross-validation for the evaluation to eliminate the weight initialisation bias and the data-split bias into training and validation sets. The metric scores of the individual models are averaged to get a robust measure independent of weight initialization and split.




Figure 8. Visualization of the Mean Distance Error (MDE) calculation between front label Q and front prediction P . The $dist(q, P)$ represents the distance between q and the pixel in P that is closest to q .

195 5.1 Evaluation Metrics

As our main measure, we use the Mean Distance Error (MDE) proposed by Gourmelon et al. (2022). It measures the distance of the predicted front \mathcal{P} to the ground truth calving front \mathcal{Q} for all images in the test set \mathcal{I} . For every pixel in the label front \mathcal{Q} , the distance to the closest pixel in the predicted front \mathcal{P} is determined. Additionally, to make the metric symmetric, the distance to the closest pixel in the label front \mathcal{Q} is determined for every pixel in the predicted front \mathcal{P} . These distances are averaged and
 200 taken as the mean distance between the two lines (see Fig. 8 and Eq. 1).

$$MDE(\mathcal{I}) = \frac{1}{\sum_{(\mathcal{P}, \mathcal{Q}) \in \mathcal{I}} (|\mathcal{P}| + |\mathcal{Q}|)} \sum_{(\mathcal{P}, \mathcal{Q}) \in \mathcal{I}} \left(\sum_{p \in \mathcal{P}} \min_{q \in \mathcal{Q}} \|p - q\|_2 + \sum_{q \in \mathcal{Q}} \min_{p \in \mathcal{P}} \|p - q\|_2 \right) \quad (1)$$

Additionally, we give classical segmentation metrics to evaluate the zone prediction. They include the Intersection over Union (IoU), which is defined as $\text{IoU} = \frac{T_p}{T_p + F_p + F_N}$, i. e., the True Positive (T_p) pixels over the sum of T_p , False Positive (F_p), and False Negative (F_N) pixels. Additionally, the F1-score is computed ($F1 = 2 \cdot \frac{pr \cdot re}{pr + re}$), which is a combination of Recall
 205 ($re = \frac{T_p}{T_p + F_N}$) and Precision ($pr = \frac{T_p}{T_p + F_p}$). Because this segmentation task is not binary, but every sample has four classes, the metrics are calculated for each class individually and then averaged with equal weighting.

5.2 Results and Discussion

In Fig. 9, the MDE of every experiment is compared. The STL approach that is trained on the front labels has an MDE of 1103 ± 72 m and the STL approach that is trained on the zone labels has an MDE of 1184 ± 225 m. A difference between
 210 the STL experiments arises in the variance of the performance, where the variance of the model trained on the zone labels is




Figure 9. Box plot of the MDE of the six experiments. Two columns are coloured equally for each experiment except for the STL experiments. The left half represents the front delineation based on the front segmentation (Task 1), and the right half represents the front delineation based on the zone segmentation (Task 2). The baseline of Gourmelon et al. (2022) is displayed as a grey line in each half.

larger. The number of samples that are incorrectly predicted to have no front pixel is similar, with an average of 27.2 ± 6.0 for training on the front and 24.8 ± 12.4 out of 122 for training with zone labels.

The lowest number of samples with false non-front-detections achieves the model that is trained on the fused labels, 3.2 ± 2.0 for the front label out of 122 samples. The baseline of Gourmelon et al. is 1 ± 1 based on the zone segmentation and 7 ± 3 from the front prediction. All values can be seen in Table A2. All MTL models have a significantly lower MDE than STL models with a significance level of $\alpha = 0.01$ using a student's t-test. The Table with the T-values of all experiment pairs is given in Table A1. The model that is additionally trained on the glacier boundary has the smallest MDE for both tasks. The MDE baseline of Gourmelon et al. is 887 ± 189 m for the front prediction and 753 ± 76 m for the zone. Overall the MDE is similar for all MTL approaches. The student's t-test shows that the differences between fused label and all other MTL approaches are insignificant ($\alpha = 0.33$).

The metrics for the zone segmentation, shown in Table A3, show a similar trend. An improvement of all MTL over the STL approach and minor changes between the MTL approaches. The STL of the nnU-Net on zone label achieves 62.4 ± 3.5 IoU and 71.7 ± 3.2 F1-score. The highest F1-score achieves the model that is trained with the additional boundary task with 81.7 ± 0.5 . The model's IoU is 72.6 ± 0.4 . The baseline of Gourmelon et al. is 80.1 ± 0.5 F1-score and 69.7 ± 0.6 IoU.

Overall, the fused label approach is the most feasible method discussed as follows. This method performs similarly well in comparison to other MTL approaches while using a minimum number of parameters and uses the nnU-Net framework with no profound changes. Moreover, this approach has a relatively small number of predictions with falsely non-detected fronts.

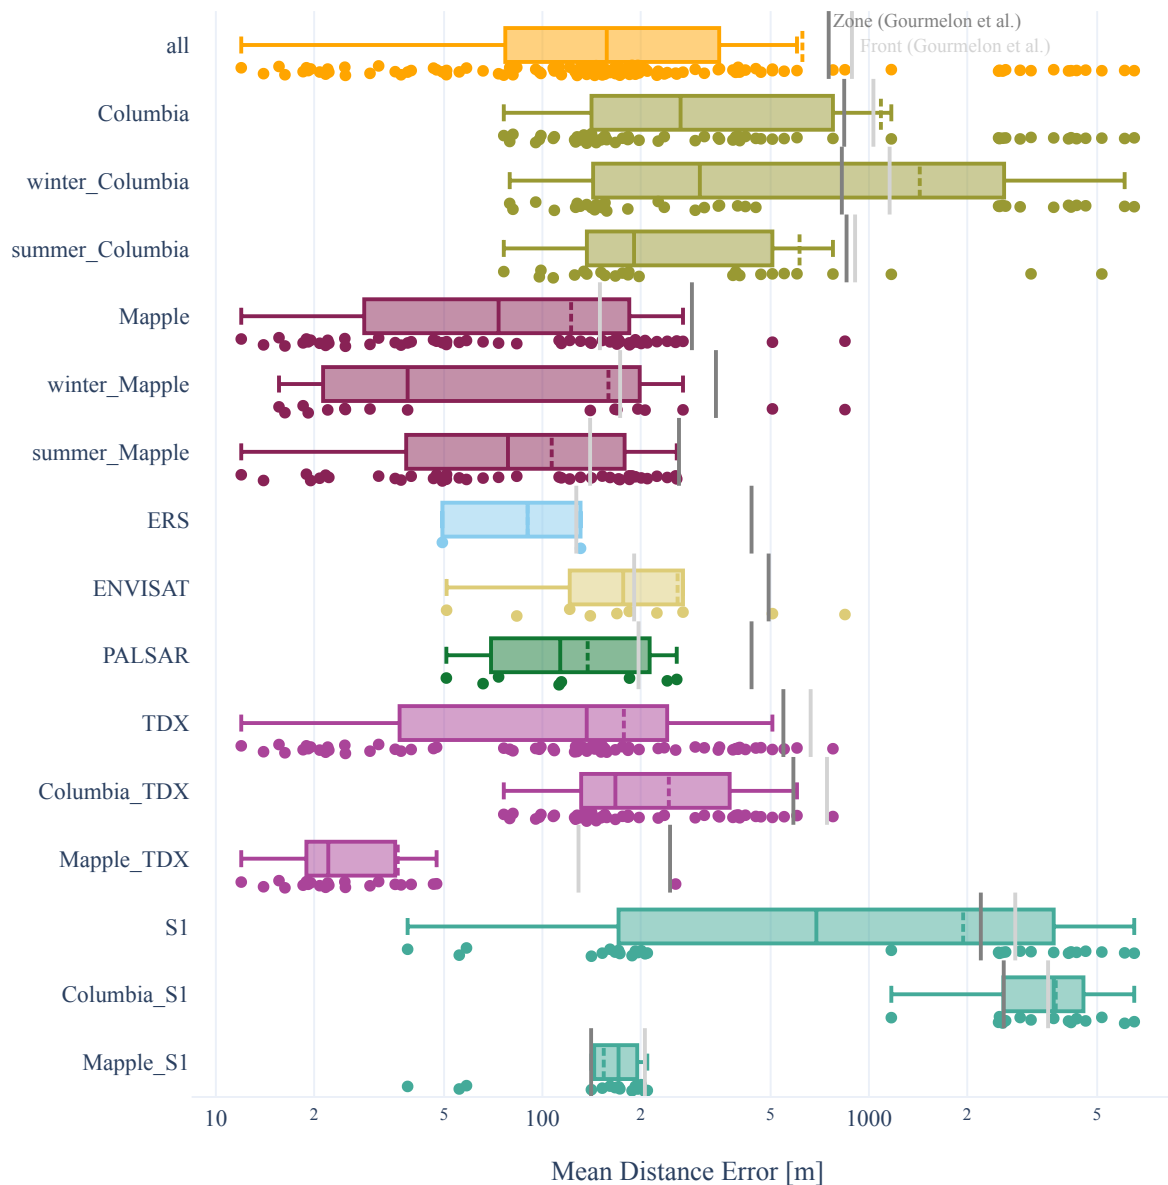


Figure 10. MDE on the test set grouped by glaciers and satellites. The box plots show the distribution of the errors. The dotted line represents the mean, and the straight line represents the median. The olive green box plots show the error on Columbia and the red samples of Mapple. The two additional rows separate the samples into winter and summer seasons. Winter in the northern hemisphere is from September to March, and winter in the southern hemisphere is from April to August. The pink box plots represent the samples taken by TDX, and the turquoise represents samples taken by Sentinel-1 (S1). The ensemble model was trained with the fused zone and front labels. The dark grey line shows the baseline of the zone prediction, and the light grey line shows the baseline of the front prediction from (Gourmelon et al., 2022). The y-axis has a logarithmic scale.



For the final evaluation, the zone segmentations of the five models that are trained on the fused labels are combined into a more robust ensemble prediction. The ensemble prediction is created by taking the mean of the probabilities for every class each model gives. The performance of the ensemble prediction results in an MDE of 515 ± 39 m and $3 \in 122$ predictions with no fronts, which is similar to the average performance of the single models. The distribution of the MDE on the test set predictions is plotted in Fig. 10. In the first row, all errors of the prediction of the 119 test samples are drawn as dots. The test set contains two glaciers: Mapple and COL. The rows below show the distribution of the MDE on the two test glaciers separated into summer and winter seasons. The main difference in MDE is between glaciers, similar to the baseline results of Gourmelon et al. (2022) with MDE of 287 ± 48 m for Mapple and 840 ± 48 m for COL. The MDE of our methods is on average 127 ± 107 m for Mapple, while the MDE of COL is 1184 ± 1761 m. The difference of the MDE is caused by a group of predictions with an error > 2000 m. The median value is 275 m for COL and 97 m for Mapple. The reasons for the large glacier differences might be related to the different shapes. Mapple has a simple calving front, a single line constraint by a straight valley. Columbia has multiple calving fronts, a stream coming from the left side of the image, one from the top, and another from the left (see Fig. 11).

There is also a seasonal difference in the MDE. The MDE of the front prediction during the summer of Mapple and COL combined have lower MDE 458 ± 1060 m than the samples captured during the winter months (996 ± 1683 m). However, the medians are closer together with 133 m in the summer months and 185 m in the winter months. The MDE baseline (Gourmelon et al., 2022) in summer is 732 ± 93 m and 776 ± 65 m in winter. Although, most outliers are from winter seasons in COL. In winter, snow coverage of the glacier is more likely, covering useful artefacts like crevasses and rock structures that can be useful for the pattern recognition of the nnU-Net. Additionally, the ocean is covered more often with ice mélange, which reduces the contrast between ocean and glacier areas.

In the lower half of Fig. 10, the MDE is grouped by satellites. The predictions for the samples captured by ERS, ENVISAT, PALSAR, and TDX have a similar average error between 150 m and 300 m. The samples captured by S1 have a higher MDE with 2060 ± 2139 m. The outliers are all from samples captured by S1 of Columbia. The neural network can not generalise from the training set to this set of samples. The 13 outliers have a front-line delineation error of > 2000 m. These are predictions from images of COL taken by S1. The low-resolution, e. g., cracks in the ice, are not visible. Fig. 11 (c) shows that the model falsely predicts ocean pixels as the front and does not predict the left calving front at all. The group of outliers heavily increases the overall MDE but can be separated clearly by specifying satellite and glacier. It shows the generalisation limits of our method.

On the other hand, the model predictions of the calving front have a high overlap with the label on images of Mapple with a resolution of ≥ 20 m (European Remote-Sensing Satellite-2 (ERS)) (see Fig. 12). There are some outliers for Mapple as well, but they are not as severe as the outliers at COL. The prediction of Mapple with a high resolution of 7 m per pixel is close to the ground truth (Fig. 12).

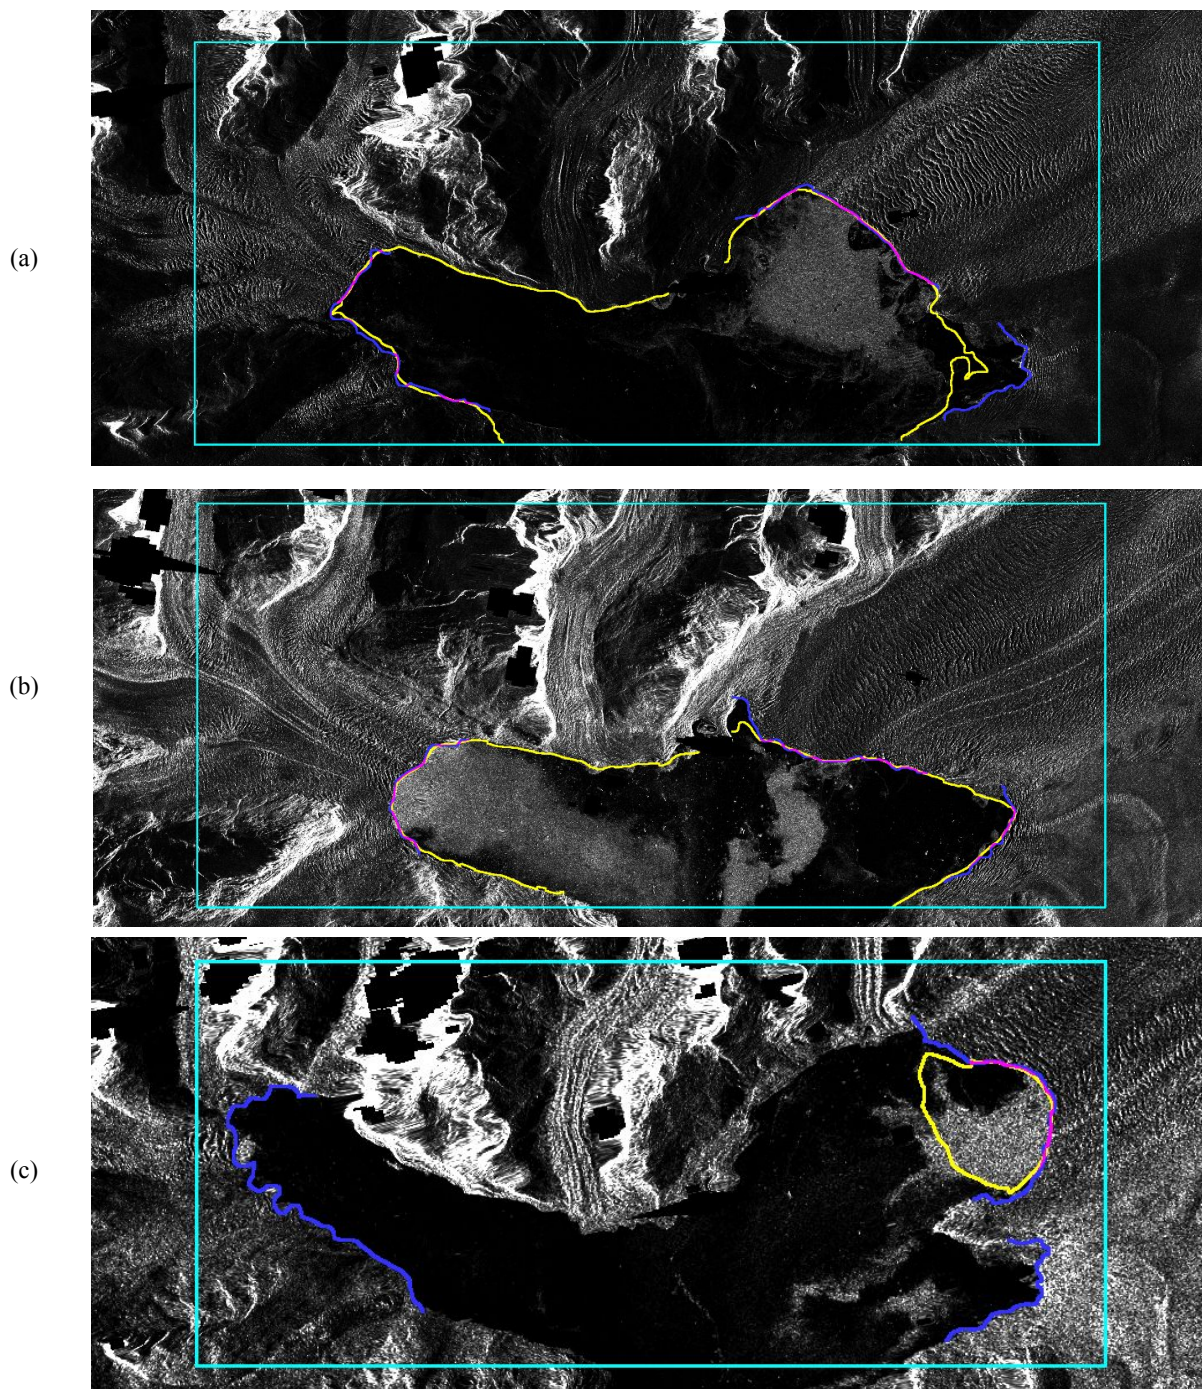


Figure 11. Calving front predictions of COL on (a) 22 June 2014, and (b) 11 February 2016 taken by TDX with 7m pixel resolution (after multi-looking with a factor of 3x3); (b) was taken by S1 on 10 September 2019 with 20 m pixel resolution; ground truth (blue), prediction (yellow), the overlap of ground truth and prediction (magenta). SAR imagery is provided by DLR, ESA, and ASF.




Figure 12. Calving front predictions of Mapple Glacier; ground truth (blue), prediction (yellow), the overlap of ground truth and prediction (magenta), bounding box (cyan). (a) was taken by TDX on 13 October 2008. (c) was taken by Phased Array type L-band SAR (PALSAR) on 29 November 2008. (c) was taken by S1 on 3 July 2016. (d) was taken by S1 on 2 March 2019. (e) was taken by ERS on 5 February 2007. (f) was taken by Environmental Satellite (ENVISAT) on 22 September 2007. SAR imagery is provided by DLR, ESA, and ASF.



6 Conclusion

260 This work explores the use of the nnU-Net by Isensee et al. (2021) to segment glacier calving fronts. The nnU-Net promises an out-of-the-box application of deep learning for segmentation tasks. We evaluate this claim on a dataset of glacier images provided by Gourmelon et al. (2022). The dataset contains two tasks. The calving front detection and the glacier zone segmentation. We try different modifications of Multi-Task-Learning with two different neural network architectures to tackle both tasks simultaneously. The results show that combining both tasks benefits each task’s performance. Still, there is no significant
 265 difference between the two network architectures. Adding more domain-specific tasks like glacier boundary delineation does not further improve the previous tasks. Due to the small area of the front line, the two labels can be fused into one label. The fusion of labels decreases the number of parameters used, shortens the training time, and reduces the deep learning expertise needed, as the nnU-Net can be used without modifications. With this approach, we achieve an average MDE of 541 ± 84 m. To improve the average MDE, future work should focus on reducing the model performance on low-resolution images.

270 The framework nnU-Net is well suited for segmenting SAR images of glaciers and calving front delineation. The modification of the nnU-Net for MTL improves the results compared to STL. However, it is not necessary for glacier segmentation and calving front delineation because both labels can be fused without losing a significant amount of information.

Appendix A: Evaluation results of all experiments

This section provides more detailed values of the evaluation. Table A1 contains the T-values for stating significant and insignificant
 275 differences between models. Table A2 contains the zone MDE and the front MDE of the six experiments. The MDE is averaged over the five model results. Additionally, the Baseline from (Gourmelon et al., 2022) is provided in the first two rows. The values correspond to Fig 9.

	Early MTL	Late MTL	Boundary MTL	Fused Labels
STL	15.78	13.46	16.14	14.29
Early MTL	0	0.31	1.02	0.43
Late MTL		0	0.95	0.58
Bound. MTL			0	0.46

(a) Front position extracted from front prediction.

	Early MTL	Late MTL	Boundary Labels	Fused Labels
STL	6.57	6.51	6.66	5.96
Early MTL	0	0.29	0.03	0.75
Late MTL		0	0.35	0.58
Bound. MTL			0	0.81

(b) Front position extracted from zone prediction.

Table A1. T-value for the significance of difference. T-values that surpass the threshold of 3.36 means that the probability for the difference to be random is below 1 %. T-values below 1 mean that the difference is by chance, with a probability of 35 %, meaning that the difference is insignificant.



Model	Experiment	Modality	MDE Front ↓ [m]	∅ Front	MDE Zone ↓ [m]	∅ Zone
Gourmelon et al.	Front	STL	887 ± 189	7 ± 3	-	-
	Zone	STL	-	-	753 ± 76	1 ± 1
nnU-Net	Front	STL	1102 ± 72	27.2 ± 6.0	-	-
	Zone	STL	-	-	1184 ± 255	24.8 ± 12.4
	Early	MTL	560 ± 25	8.6 ± 3.0	509 ± 43	2.2 ± 0.4
	Late	MTL	568 ± 51	8.4 ± 3.5	516 ± 40	4.0 ± 1.1
	Boundary	MTL	543 ± 27	5.6 ± 1.4	509 ± 19	4.4 ± 1.8
	Fused	MTL	552 ± 33	3.2 ± 2.0	541 ± 84	3.4 ± 1.3

Table A2. MDE of all six experiments. Every experiment setup is trained five times with different weight initialization and train-validation-split and then averaged. The columns show the number of images for which falsely no front was detected out of the 122 test samples.

Model	Experiment	Modality	Precision↑	Recall↑	F1↑	IoU↑
Gourmelon et al.	Zone	STL	84.2 ± 0.5	79.6 ± 0.9	80.1 ± 0.5	69.7 ± 0.6
nnU-Net	Front	STL	-	-	-	-
	Zone	STL	81.2 ± 2.4	71.7 ± 3.2	71.9 ± 3.7	62.4 ± 3.5
	Early	MTL	87.4 ± 0.4	80.9 ± 0.7	81.6 ± 0.6	72.6 ± 0.9
	Late	MTL	86.4 ± 0.4	79.9 ± 0.6	80.7 ± 0.7	71.1 ± 0.7
	Boundary	MTL	87.5 ± 0.3	80.8 ± 0.4	81.7 ± 0.5	72.6 ± 0.4
	Fused	MTL	87.0 ± 0.2	79.1 ± 1.7	80.7 ± 0.1	70.8 ± 1.8

Table A3. Segmentation metrics of all six experiments. Every experiment setup is trained five times with different weight initialization and train-validation-split. The metric results of each run are then averaged.

Appendix B: Calving front labels

The other six glacier sites in this section are displayed complementary to Fig 4. The corresponding calving fronts of the different timesteps are coloured with a gradient from bright for the past and dark red for the most recent fronts.

Appendix C: Temporal distribution of zone label

Figure C1 show the temporal distributio of the zone label. Jakobshavn Isbrae (JAK) and COL, the two glaciers in the northern hemisphere, show an increase in ocean area with a decreasing glacier area. In particular, the set of images of the glaciers on the AP has samples with large areas with no available information. Low partial coverage by the radar swath causes prominent peaks of the no information available class. JAK's area distribution shows a repetitive structure of increasing and decreasing glacier area. This pattern represents seasonal changes.

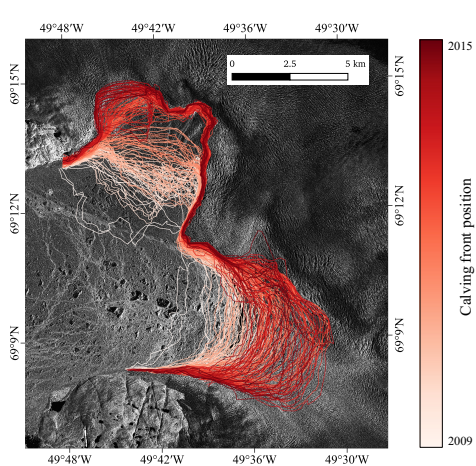


Figure B1. Jakobshavn Isbrae Glacier from 1995 to 2014.




Figure B2. Mapple Glacier from 2006 to 2020.

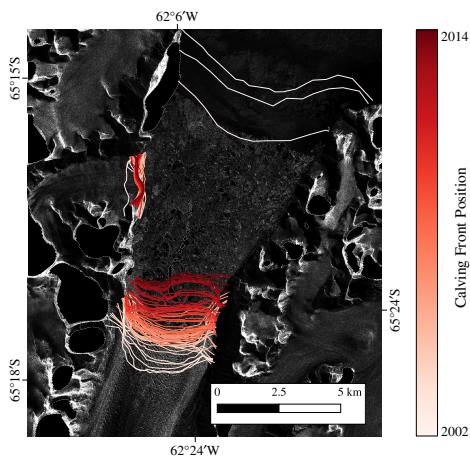


Figure B3. Crane Glacier from 2002 to 2014.




Figure B4. DBE Glacier from 1995 to 2014.




Figure B5. Jorum Glacier from 2003 to 2020.




Figure B6. Sjögren-Inlet Glacier from 1995 to 2014.




Figure C1. Changes of the class distributions in the zone labels of each glacier over time. The ocean is coloured in blue, the glacier area in white, the rock in grey and the area with no available information (NA) in black. There is no radar reflection in the NA area due to terrain elevation causing shadows or due to limited coverage by the radar swath.



Code and data availability

The code is fundamentally based on the nnU-Net by Isensee et al. (2021), which is also publicly available (<https://github.com/MIC-DKFZ/nnUNet>). We modified the neural network architecture and corresponding pre- and post-processing. The modified version of the nnU-Net and our evaluation and visualisation scripts is available (https://gitlab.cs.fau.de/ho11laqe/nnunet_glacer.git).
290 The dataset is provided by Gourmelon et al. (2022) (<https://doi.org/10.1594/PANGAEA.940950>).

Author contribution

OH was responsible for conceptualizing the study and developing the methodology and software for the experiments. NG and TS provided the dataset and a benchmark score. NG and VC supervised the study and reviewed and edited the manuscript. OH
295 prepared the manuscript with contributions from all co-authors. JF, VC, and NG reviewed and edited the manuscript.

Competing interests

The authors declare that they have no conflict of interest.

Acknowledgments

J. J. Fürst has received funding from the European Union's Horizon 2020 research and innovation programme via the European Research Council (ERC) as a Starting Grant (StG) under grant agreement No. 948290. T. Seehaus was supported
300 by the ESA Living Planet Fellowship project "MIT-AP". We acknowledge the free provision of the SAR data via various proposals from ESA, ASF, and DLR. The authors gratefully acknowledge the scientific support and HPC resources provided by the Erlangen National High-Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). The hardware is funded by the German Research Foundation (DFG). We acknowledge financial support by Deutsche Forschungsgemeinschaft and Friedrich-Alexander-Universität Erlangen-Nürnberg within the funding
305 programme "Open Access Publication Funding". This research has been supported by the Friedrich-Alexander-Universität Erlangen-Nürnberg (Emerging Fields Initiative "Tapping the Potential of Earth Observation") and the STAEDTLER Foundation. The authors thank the Elitenetzwerk Bayern for financial support under the International Doctorate Program "Measuring and modelling mountain glaciers and ice caps in a changing climate (M^3 OCCA)".



310 References

- Abolvardi, A. A., Hamey, L., and Ho-Shon, K.: UNET-Based Multi-Task Architecture for Brain Lesion Segmentation, in: *Digital Image Computing: Techniques and Applications (DICTA)*, pp. 1–7, <https://doi.org/10.1109/DICTA51227.2020.9363397>, 2020.
- Amundson, J. M., Fahnestock, M., Truffer, M., Brown, J., Lüthi, M. P., and Motyka, R. J.: Ice mélange dynamics and implications for terminus stability, Jakobshavn Isbrse, Greenland, *Journal of Geophysical Research: Earth Surface*, 115, <https://doi.org/10.1029/2009JF001405>,
315 2010.
- Amyar, A., Modzelewski, R., Li, H., and Ruan, S.: Multi-task deep learning based CT imaging analysis for COVID-19 pneumonia: Classification and segmentation, *Computers in Biology and Medicine*, 126, 104 037, <https://doi.org/10.1016/j.compbimed.2020.104037>, 2020.
- Baumhoer, C. A., Dietz, A. J., Dech, S., and Kuenzer, C.: Remote sensing of antarctic glacier and ice-shelf front dynamics-a review, *Remote Sensing*, 10, 1445, <https://doi.org/10.3390/rs10091445>, 2018.
- 320 Baumhoer, C. A., Dietz, A. J., Kneisel, C., and Kuenzer, C.: Automated Extraction of Antarctic Glacier and Ice Shelf Fronts from Sentinel-1 Imagery Using Deep Learning, *Remote Sensing*, 11, 2529, <https://doi.org/10.3390/rs11212529>, 2019.
- Beer, C., Biebow, N., Braun, M., Döring, N., Gaedicke, C., Gutt, J., Hagen, W., Hauck, J., Heinemann, G., Herata, H., et al.: *Forschungsagenda Polarregionen im Wandel*, p. 79, Bundesministerium für Bildung und Forschung (BMBF), Germany, 2021.
- Bischke, B., Helber, P., Folz, J., Borth, D., and Dengel, A.: Multi-Task Learning for Segmentation of Building Foot-
325 prints with Deep Neural Networks, in: *International Conference on Image Processing (ICIP)*, pp. 1480–1484, IEEE, Taipei, <https://doi.org/10.1109/ICIP.2019.8803050>, 2019.
- Carr, J. R., Stokes, C., and Vieli, A.: Recent retreat of major outlet glaciers on Novaya Zemlya, Russian Arctic, influenced by fjord geometry and sea-ice conditions, *Journal of Glaciology*, 60, 155–170, <https://doi.org/10.3189/2014JoG13J122>, 2014.
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H.: Encoder-Decoder with Atrous Separable Convolution for Semantic Image
330 Segmentation, in: *European conference on computer vision (ECCV)*, edited by Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y., pp. 833–851, Springer International Publishing, Cham, https://doi.org/https://doi.org/10.1007/978-3-030-01234-2_49, 2018.
- Chen, S., Bortsova, G., García-Uceda Juárez, A., van Tulder, G., and de Bruijne, M.: Multi-task Attention-Based Semi-supervised Learning for Medical Image Segmentation, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, edited by Shen, D., Liu, T., Peters, T. M., Staib, L. H., Essert, C., Zhou, S., Yap, P.-T., and Khan, A., *Lecture Notes in Computer Science*, pp. 457–465,
335 Springer International Publishing, Cham, https://doi.org/10.1007/978-3-030-32248-9_51, 2019.
- Cheng, D., Hayes, W., Larour, E., Mohajerani, Y., Wood, M., Velicogna, I., and Rignot, E.: Calving Front Machine (CALFIN): glacial termini dataset and automated deep learning extraction method for Greenland, 1972–2019, *The Cryosphere*, 15, 1663–1675, <https://doi.org/10.5194/tc-15-1663-2021>, 2021.
- Davari, A., Baller, C., Seehaus, T., Braun, M., Maier, A., and Christlein, V.: Pixel-wise Distance Regression for Glacier Calving Front Detection and Segmentation, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–10, <https://doi.org/10.1109/TGRS.2022.3158591>,
340 2022.
- Friedl, P., Seehaus, T. C., Wendt, A., Braun, M. H., and Höppner, K.: Recent dynamic changes on Fleming Glacier after the disintegration of Wordie Ice Shelf, Antarctic Peninsula, *The Cryosphere*, 12, 1347–1365, <https://doi.org/10.5194/tc-12-1347-2018>, 2018.
- Gourmelon, N., Seehaus, T., Braun, M., Maier, A., and Christlein, V.: Calving fronts and where to find them: a benchmark dataset and methodology for automatic glacier calving front extraction from synthetic aperture radar imagery, *Earth System Science Data*, 14, 4287–4313, <https://doi.org/10.5194/essd-14-4287-2022>, 2022.



- Gourmelon, N., Seehaus, T., Braun, M. H., Maier, A., and Christlein, V.: CaFFe (CALving Fronts and where to Find thEm: a benchmark dataset and methodology for automatic glacier calving front extraction from sar imagery), <https://doi.org/10.1594/PANGAEA.940950>, 2022.
- 350 Hartmann, A., Davari, A., Seehaus, T., Braun, M., Maier, A., and Christlein, V.: Bayesian U-Net for Segmenting Glaciers in Sar Imagery, *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 41, 3479–3482, <https://doi.org/10.1109/IGARSS47720.2021.9554292>, ISSN: 2153-7003, 2021.
- He, K., Lian, C., Zhang, B., Zhang, X., Cao, X., Nie, D., Gao, Y., Zhang, J., and Shen, D.: HF-UNet: Learning Hierarchically Inter-Task Relevance in Multi-Task U-Net for Accurate Prostate Segmentation in CT Images, *IEEE transactions on medical imaging*, 40, 2118–2128, 355 <https://doi.org/10.1109/TMI.2021.3072956>, 2021.
- Heidler, K., Mou, L., Baumhoer, C., Dietz, A., and Zhu, X. X.: HED-UNet: Combined Segmentation and Edge Detection for Monitoring the Antarctic Coastline, *IEEE Transactions on Geoscience and Remote Sensing*, pp. 1–14, <https://doi.org/10.1109/TGRS.2021.3064606>, 2021.
- Heller, N., Isensee, F., Maier-Hein, K. H., Hou, X., Xie, C., Li, F., Nan, Y., Mu, G., Lin, Z., Han, M., Yao, G., Gao, Y., Zhang, Y., 360 Wang, Y., Hou, F., Yang, J., Xiong, G., Tian, J., Zhong, C., Ma, J., Rickman, J., Dean, J., Stai, B., Tejpaul, R., Oestreich, M., Blake, P., Kaluzniak, H., Raza, S., Rosenberg, J., Moore, K., Walczak, E., Rengel, Z., Edgerton, Z., Vasdev, R., Peterson, M., McSweeney, S., Peterson, S., Kalapara, A., Sathianathen, N., Papanikolopoulos, N., and Weight, C.: The state of the art in kidney and kidney tumor segmentation in contrast-enhanced CT imaging: Results of the KiTS19 challenge, *Medical Image Analysis*, 67, 101 821, <https://doi.org/https://doi.org/10.1016/j.media.2020.101821>, 2021.
- 365 Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., and Maier-Hein, K. H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation, *Nature Methods*, 18, 203–211, <https://doi.org/10.1038/s41592-020-01008-z>, 2021.
- Jang, H.-J. and Cho, K.-O.: Applications of deep learning for the analysis of medical data, *Archives of pharmacal research*, 42, 492–504, <https://doi.org/https://doi.org/10.1007/s12272-019-01162-9>, 2019.
- Kholiavchenko, M., Sirazitdinov, I., Kubrak, K., Badrutdinova, R., Kuleev, R., Yuan, Y., Vrtovec, T., and Ibragimov, B.: Contour-aware multi-label chest X-ray organ segmentation, *International Journal of Computer Assisted Radiology and Surgery*, 15, 425–436, 370 <https://doi.org/10.1007/s11548-019-02115-9>, 2020.
- Kneib-Walter, A., Lüthi, M. P., Moreau, L., and Vieli, A.: Drivers of Recurring Seasonal Cycle of Glacier Calving Styles and Patterns, *Frontiers in Earth Science*, 9, <https://doi.org/10.3389/feart.2021.667717>, 2021.
- Li, X., Wang, Y., Tang, Q., Fan, Z., and Yu, J.: Dual U-Net for the Segmentation of Overlapping Glioma Nuclei, *IEEE Access*, 7, 84 040– 375 84 052, <https://doi.org/10.1109/ACCESS.2019.2924744>, 2019.
- Loebel, E., Scheinert, M., Horwath, M., Heidler, K., Christmann, J., Phan, L. D., Humbert, A., and Zhu, X. X.: Extracting glacier calving fronts by deep learning: the benefit of multi-spectral, topographic and textural input features, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–12, <https://doi.org/10.1109/TGRS.2022.3208454>, 2022.
- McNabb, R. W., Hock, R., and Huss, M.: Variations in Alaska tidewater glacier frontal ablation, 1985–2013, *Journal of Geophysical Research: Earth Surface*, 120, 120–136, <https://doi.org/https://doi.org/10.1002/2014JF003276>, 2015. 380
- Minowa, M., Schaefer, M., Sugiyama, S., Sakakibara, D., and Skvarca, P.: Frontal ablation and mass loss of the Patagonian icefields, *Earth and Planetary Science Letters*, 561, 116 811, <https://doi.org/https://doi.org/10.1016/j.epsl.2021.116811>, 2021.
- Mohajerani, Y., Wood, M., Velicogna, I., and Rignot, E.: Detection of Glacier Calving Margins with Convolutional Neural Networks: A Case Study, *Remote Sensing*, 11, 74, <https://doi.org/10.3390/rs11010074>, 2019.



- 385 Periyasamy, M., Davari, A., Seehaus, T., Braun, M., Maier, A., and Christlein, V.: How to Get the Most Out of U-Net for Glacier Calving Front Segmentation, *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15, 1712–1723, <https://doi.org/10.1109/JSTARS.2022.3148033>, 2022.
- Recinos, B., Maussion, F., Rothenpieler, T., and Marzeion, B.: Impact of frontal ablation on the ice thickness estimation of marine-terminating glaciers in Alaska, *The Cryosphere*, 13, 2657–2672, <https://doi.org/10.5194/tc-13-2657-2019>, 2019.
- 390 Recinos, B., Maussion, F., Noël, B., Möller, M., and Marzeion, B.: Calibration of a frontal ablation parameterisation applied to Greenland’s peripheral calving glaciers, *Journal of Glaciology*, 67, 1177–1189, <https://doi.org/10.1017/jog.2021.63>, 2021.
- Robel, A. A., Schoof, C., and Tziperman, E.: Persistence and variability of ice-stream grounding lines on retrograde bed slopes, *The Cryosphere*, 10, 1883–1896, <https://doi.org/10.5194/tc-10-1883-2016>, 2016.
- Ronneberger, O., Fischer, P., and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, edited by Navab, N., Hornegger, J., Wells, W. M., and Frangi, A. F., vol. 9351, pp. 234–241, Springer International Publishing, Cham, https://doi.org/10.1007/978-3-319-24574-4_28, 2015.
- 395 Rott, H., Wuite, J., Rydt, J. D., Gudmundsson, G. H., Floricioiu, D., and Rack, W.: Impact of marine processes on flow dynamics of northern Antarctic Peninsula outlet glaciers, *Nature Communications*, 11, 2969, <https://doi.org/10.1038/s41467-020-16658-y>, 2020.
- Shepherd, A., Ivins, E., Rignot, E., Smith, B., Broeke, M. V. D., Velicogna, I., Whitehouse, P., Briggs, K., Joughin, I., Krinner, G., Nowicki, 400 S., Payne, T., Scambos, T., Schlegel, N., Geruo, A., Agosta, C., Ahlstrøm, A., Babonis, G., Barletta, V., Blazquez, A., Bonin, J., Csatho, B., Cullather, R., Felikson, D., Fettweis, X., Forsberg, R., Gallee, H., Gardner, A., Gilbert, L., Groh, A., Gunter, B., Hanna, E., Harig, C., Helm, V., Horvath, A., Horwath, M., Khan, S., Kjeldsen, K. K., Konrad, H., Langen, P., Lecavalier, B., Loomis, B., Luthcke, S., McMillan, M., Melini, D., Mernild, S., Mohajerani, Y., Moore, P., Mouginot, J., Moyano, G., Muir, A., Nagler, T., Nield, G., Nilsson, J., Noel, B., Ootaka, I., Pattle, M. E., Peltier, W. R., Pie, N., Rietbroek, R., Rott, H., Sandberg-Sørensen, L., Sasgen, I., Save, H., Scheuchl, 405 B., Schrama, E., Schröder, L., Seo, K. W., Simonsen, S., Slater, T., Spada, G., Sutterley, T., Talpe, M., Tarasov, L., Berg, W. J. V. D., Wal, W. V. D., Wessem, M. V., Vishwakarma, B. D., Wiese, D., and Wouters, B.: Mass balance of the Antarctic Ice Sheet from 1992 to 2017, *Nature*, 558, 219–222, <https://doi.org/10.1038/s41586-018-0179-y>, 2018.
- Straneo, F., Heimbach, P., Sergienko, O., Hamilton, G., Catania, G., Griffies, S., Hallberg, R., Jenkins, A., Joughin, I., Motyka, R., Pfeffer, W. T., Price, S. F., Rignot, E., Scambos, T., Truffer, M., and Vieli, A.: Challenges to Understanding the Dynamic Response of Greenland’s 410 Marine Terminating Glaciers to Oceanic and Atmospheric Forcing, *Bulletin of the American Meteorological Society*, 94, 1131–1144, <https://doi.org/10.1175/BAMS-D-12-00100.1>, 2013.
- Zhang, E., Liu, L., and Huang, L.: Automatically delineating the calving front of Jakobshavn Isbræ from multitemporal TerraSAR-X images: a deep learning approach, *The Cryosphere*, 13, 1729–1741, <https://doi.org/10.5194/tc-13-1729-2019>, 2019.
- Zhang, E., Liu, L., Huang, L., and Ng, K. S.: An automated, generalized, deep-learning-based method for delineating the calving 415 fronts of Greenland glaciers from multi-sensor remote sensing imagery, *Remote Sensing of Environment*, 254, 112265, <https://doi.org/10.1016/j.rse.2020.112265>, 2021.