

## Response to the second round of revisions on "Out-of-the-box calving front detection method using deep learning" by Oskar Herrmann et al.

First of all, we want to thank the reviewers again for their constructive comments on our revised manuscript. The following pages contain a list of reviewer comments followed by our replies. The comments are sequentially numbered and associated with the corresponding reviewer. The document contains active cross-references between the replies and the modifications in the main text. The black label references the original reviewer's comment within the manuscript. For instance, the reference **1.3**, refers to the response from the third comment of the first reviewer, the reference **A2/C2/D2 (3.5)**, which is typeset in the manuscript margin, refers to the second addition/change/deletion stemming from the fifth comment of the third reviewer.

### Comments of the 1st Reviewer:

**1.1** *My main concern is about the impact of this study. The ultimate goal of automating terminus extraction is to provide as many as quality-controlled termini for glaciological community to conduct scientific research. I understand it might not be the scope of this study to achieve this such an ultimate goal, but I concern that a high uncertainty might undermine the impact of this paper, even though this is a methodology paper. I believe the key contribution of this study is to improve the error from 753m to 541 m, compared with Gourmelon et al. (2022). However, both errors are relatively high.*

- ▶ Thank you for your comment. We invite you to look at the animations uploaded to Zenodo (<https://zenodo.org/record/8379954>). We provided an animation of the two test sites. The Mean Distance Error (MDE) seems to be high, but the animation gives a better insight into the weaknesses and strengths of the method. You can see that in the images of the Columbia Glacier, the predicted calving front is longer than the labelled front and extends on the coastline. This has a negative influence on the MDE, but the predictions are still useful to show changes in the calving front. If the coastline is known, the misclassified part can easily be removed. In nearly all predictions of Columbia (COL) taken by TanDEM-X (TDX) the northern coastline was misclassified as a calving front. It is a difficult task to distinguish this particular coastline from the calving front because the glacier retreated so far that it transitioned from a marine-terminating glacier to a land-terminating glacier. Since the glacier front is close to the coastline, even humans find it difficult to distinguish between the two.

The median MDE is another representation of the performance of the method and gives a different impression. The median is the value that separates the lower and the higher half of the set of MDE of the test set samples. In our case, this is 153 m, which is also shown in Fig. 10. This shows that half of the predictions have an MDE below 153 m. The mean MDE is heavily influenced by a group of outliers, with MDE greater than 2000 m in contrast to the median MDE. If you would exclude these twelve outliers, the mean MDE would drop from 541 m to 181 m. Similar post-processing steps were used by Zhang et al. (2021), where a complexity score is computed for each front. Predictions above a certain threshold were discarded, and additionally, predictions were discarded manually. In another work from Heidler et al. (2021) the metric is only calculated within a 2000 m margin of the true coastline. Davari et al. (2022b) use a tolerance belt around the ground truth.

We dispense with such techniques to reduce the MDE to maintain comparability to the baseline from Gourmelon et al. (2022) and to show the "out-of-the-box" performance of the network itself. If the model is used for calving front detection, we can recommend the aforementioned post-processing steps using the position of the coastline.

We added another point for discussion in addition 3 and updated the Zenodo link in the manuscript (changes 2 and 3).

**1.2** *2.2 In Line 66, the author mentioned that the nn-U-Net is a framework for training the U-Net and can adjust the hyperparameter accordingly. Therefore, in the previous comments, I was wondering which hyperparameters the framework adjust and how the it is done based on the dataset automatically (accordingly). In the response, it seems to me that the framework did not adjust the hyperparameter according to the dataset. The patch size and the batch size are decided by author based on the GPU memory, which is commonly adopted when using any other deep learning methods. So, I suggest to remove the sentence about the nn-UNet here can adjust the hyperparameters according to the dataset. With that being said, I am still wondering that why the nn-UNet here does not adjust the hyperparameters, as being a self-configuring method is the most important feature of the nn-UNet (Isensee et al., 2021).*

- ▶ Thank you for pointing out the confusing sentence. Our intention in line 66 was to introduce the nnU-Net in general. As you correctly point out, the self-adjusting hyperparameters play a minor role in our work, because most rule-based hyperparameters

are only used for three-dimensional (3D) segmentation with different spacings between pixels. Therefore, we rephrased the sentence in change 1 to emphasize the fixed hyperparameter over the rule-based hyperparameters.

We mentioned that we set the patch and batch size to maintain comparability between our experiments, but the initial value is generated by the nnU-Net with the early-branching architecture based on the median size of the training samples and the GPU memory. We slightly changed section 4.1, to make this clearer (see deletion 1 and addition 1).

**1.3** *2.3 If “out-of-the-box” means “unmodified”, why it is important to highlight in the title that the author use the unmodified nn-UNet? In line 72, the second contribution is the out-of-the-box application of the nn-UNet. It seems a little odd to me to highlight “unmodified” application of a method.*

- ▶ We want to emphasise the “unmodified” use of a method in the title because we demonstrate the generalisation capabilities of the nnU-Net from the medical domain to a completely unrelated domain of “remote sensing”. In addition, we modified the nnU-Net for Multi-Task-Learning (MTL) and showed that the “unmodified” architecture worked equally well with fused “modified” labels. Referring to line 72, the contribution of applying an unmodified framework to a new dataset lies in (1) installing the nnU-net, (2) converting the dataset into pseudo-3D Neuroimaging Informatics Technology Initiative (NIfTI), which is a file format used for Computer Tomography (CT) and Magnetic Resonance Imaging (MRI) scans, (3) training the nnU-Net on five folds of the dataset, which took five days on an Nvidia RTX3080, for each set of labels (fronts and zones) (4) converting the output back from pseudo-3D to further usable Portable Network Graphic (PNG) and (5) evaluating the results using MDE (see table A2) and Segmentation metrics Precision, Recall, F1-score and Intersection over Union (IoU) (see table A3). The files for the data conversion can be found in our GitHub project in the folder nnunet/dataset\_conversion in the files Task501\_Glacier\_front.py, Task501\_Glacier\_reverse.py, Task502\_Glacier\_zone.py, Task502\_Glacier\_reverse.py. We don’t say that this is our main contribution, but it is at least worth mentioning. From our perspective, the main contribution are points three, four and five from line 74, which refer to the implementation and evaluation of different ways to include both labels (zone and front) into the same training of the nnU-Net. The title refers to our conclusion of the comparison: all experiments that incorporate both labels perform better than the single-task experiments, emphasizing that the out-of-the-box architecture trained on the fused labels performs as well as the invasive experiments with a different architecture. We rephrased the contributions in section 1.

**1.4** *I tried HuggingFace link, however it shows error after I upload one image of a glacier and click submit. I would suggest to add additional information (e.g., image format) in the webpage or the manuscript to make the link work.*

- ▶ Thank you for trying it out. We are sorry that it didn’t work for you. This can have multiple reasons: (1) The input image was not a PNG. Since the network was trained on PNG of the dataset, we stick to this file format. (2) The image was too large. Unfortunately, this platform provides limited computational resources, we recommend a maximum size of 200KB, which takes about 50 seconds to process. For larger files the HuggingFace project can be downloaded and run on a local PC. We elaborated the description of the interface and included two examples (see [https://huggingface.co/spaces/ho11laqe/nnUNet\\_calvingfront\\_detection](https://huggingface.co/spaces/ho11laqe/nnUNet_calvingfront_detection)).

**1.5** *2.13 If I understand correctly, in the multi-task learning experiments, the task of delineating calving front only and delineating the whole glacier boundary are abandoned in the inference, but only help to adjust the parameters in the training progress.*

- ▶ Thank you for highlighting this fact. Yes, during inference the other prediction is omitted. We made it a bit clearer in addition 2.

**1.6** *Another related issue is about the fused label. If the author fuse the label of calving front into the zone label with an additional class, I think it is a single classification task instead of a multi-task method. It is the same with the original zone segmentation but with one more class representing the calving front. Also, as this fused label yields the best results, it might be better to change the conclusion throughout the manuscript regarding the multi-task-learning.*

- ▶ Thank you for your comment. This is a point of discussion. From a technical point of view, we fuse two different labels into one and we can use the unmodified Single-Task-Learning (STL) architecture. Nevertheless, they come from two different tasks,

and as such we refer to it as MTL since we use both labels in contrast to the true STL experiments. We show this ambiguity in Fig. 6. We try to circumvent the hard categorisation of the fused label training into STL or MTL, because of the ambiguity. The conclusion still holds true as we state that "The results show that combining both tasks increases each task's performance.". We rephrased another sentence in addition 4,

### **Comments of the 2nd Reviewer:**

#### **2.1** *Fig 3e: For some reason the front of Jacobshavn is almost not visible.*

- ▶ Thank you for pointing this out. The different resolution and size of the images is the reason for the calving front to appear thinner. We dilated the front label in Fig. 3, for better visibility.

#### **2.2** *L 146: Would it be possible to approach Gourmelon et al. to get information about the polarizations? The person that pre-processed the images should now the details about the SAR data.*

- ▶ Thank you for your request. We approached the person, but unfortunately, it was not possible to receive information about the polarisation. The polarization was not logged explicitly. Retroactively researching the polarization would be out of the scope of this work.

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

Oskar Herrmann<sup>1</sup>, Nora Gourmelon<sup>2</sup>, Thorsten Seehaus<sup>1</sup>, Andreas Maier<sup>2</sup>, Johannes J. Fürst<sup>1</sup>, Matthias H. Braun<sup>1</sup>, and Vincent Christlein<sup>2</sup>

<sup>1</sup>Institute of Geography, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

<sup>2</sup>Pattern 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  $753 \pm 76$  m to  $541 \pm 84$  m. The scripts for our experiments are published on GitHub ([https://github.com/ho11laqe/nnUNet\\_calvingfront\\_detection](https://github.com/ho11laqe/nnUNet_calvingfront_detection)). An easy access version is published on Hugging Face ([https://huggingface.co/spaces/ho11laqe/nnUNet\\_calvingfront\\_detection](https://huggingface.co/spaces/ho11laqe/nnUNet_calvingfront_detection)).

## 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; Mohajerani et al., 2019; Baumhoer et al., 2021; Zhang et al., 2021; Heidler et al., 2021; Marochov et al., 2020; Loebel et al., 2022). Calving fronts can be located in both optical and SAR imagery. In optical imagery, calving fronts are more easily distinguishable, whereas SAR imagery has a higher scene availability, as it is independent of daytime, season and cloud coverage (Baumhoer et al., 2018). A direct comparison between the results of existing deep learning-based calving front extraction studies is not possible, as the models have been trained on different data, tested on different test sets, and evaluated using slightly differing metrics.

The benchmark 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  $67.7 \pm 0.6$ . 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  $753 \pm 76$  m. The model trained directly on the front label achieved an MDE of  $887 \pm 189$  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).

65 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. The framework contains proven deep learning techniques and established hyperparameter values, as well as rule-based parameters that depend on properties of the dataset and available GPU-memory.~~ 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 baseline for our evaluation of the nnU-Net, we use the results of Gourmelon et al. (2022). <sup>CI (1,2)</sup>

In particular, our contributions are as follows:

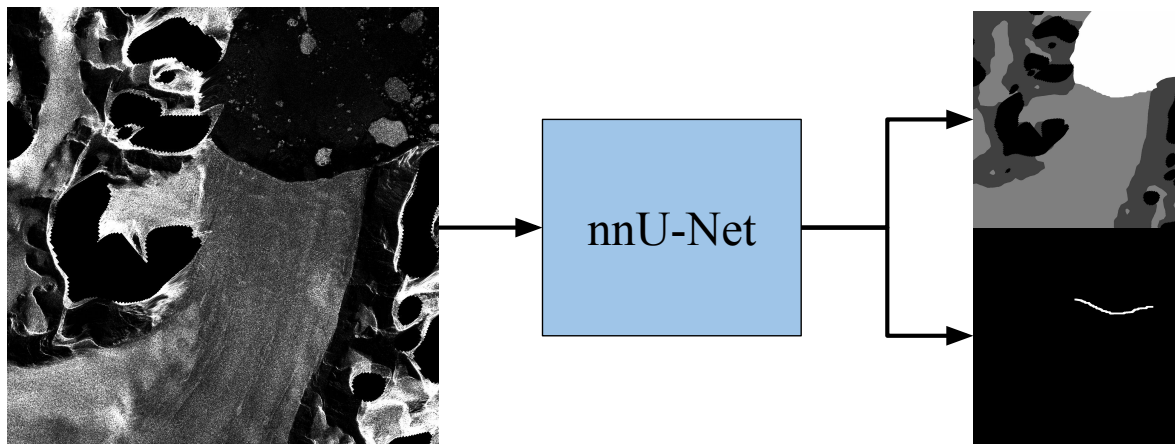
(1) Visualizations showing the temporal and spatial distribution of the dataset by Gourmelon et al. (2022). (2) Application and evaluation of the nnU-Net for calving front detection and zone segmentation. (3) Two different modifications of the original nnU-Net to incorporate both labels for multi-task learning. (4) Test if an artificial third label improves the calving front detection (5) Introduction of an efficient approach that fuses the two labels of the dataset. (6) Analysis of the influence of season, glacier, and 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

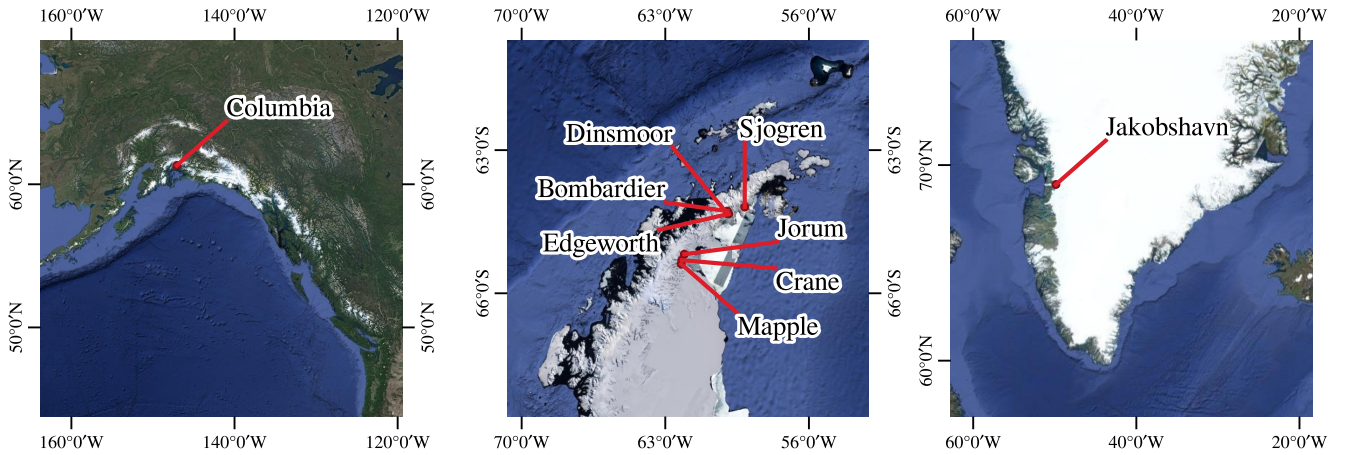
Automated monitoring of glacier-covered areas is a growing research field. Recent glacier monitoring uses deep learning methods due to the increasing availability of satellite images and computing power. Many methods are based on the U-Net (Ronneberger et al., 2015). They modify the vanilla U-Net for better performance on calving front detection (Loebel et al., 2022; Mohajerani et al., 2019; Zhang et al., 2019). 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., 2022a). 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).



**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.

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 from the reference of 345 m compared to 483 m with the vanilla U-Net. To avoid distortions of the metric from areas far from the coast, the metric is calculated within 2 km of the true coastline. A further development of the HED-UNet for Antarctic ice shelf front detection is proposed by Baumhoer et al. (2023). Their post-processing includes an elevation threshold of 110 m to remove erroneous classifications in the high-altitude dry snow zones.



**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

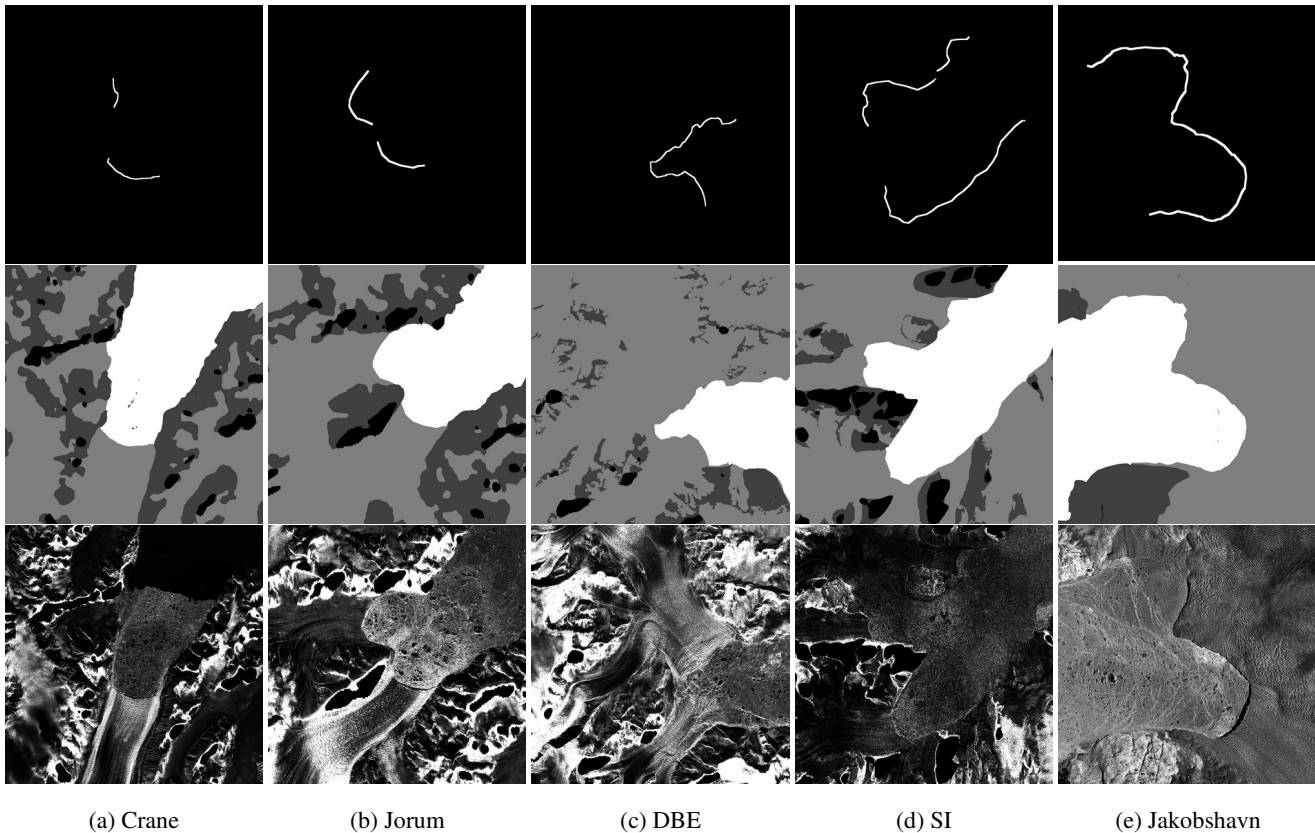
**Table 1.** Properties of the dataset, including a list of captured glaciers, train-test split, number of images per glacier, and covered area.

### 110 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 (COL) Glacier in Alaska and Jakobshavn Isbrae (JAK) (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  
 115 closest to the south pole, followed by Dinsmoore, Bombardier, Edgeworth (DBE). They are located so close together that they are treated as one glacier site. The last glacier is the Sjøgren-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  
 120 in the dataset. The other glaciers are outlet glaciers of polar ice sheets. Therefore, Columbia can be seen as a benchmark glacier for the generalization capability of the segmentation method. Columbia and Mapple also differ in the shape of the glacier. The






**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). Black dots in the ocean zone of Crane are areas with no data available due to a small artefact at the former calving front of Crane Glacier in the DEM (Cook et al., 2012) used for the orthorectification of the SAR data. SAR imagery is provided by DLR, ESA, and ASF.

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.

125     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




**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.

130 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 calving fronts, and the dark red lines the more recent positions. The constant retreat of COL is visible through the tiered lines. The  
 135 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. The resolution refers to the ground range resolution. Environmental Satellite (ENVISAT), European Remote-Sensing Satellite-2 (ERS), Sentinel-1 (S1) has  
 140 a resolution of 20 m, Phased Array type L-band SAR (PALSAR) has a 17 m resolution, and TDX has a 7 m per pixel resolution. Unfortunately, the dataset does not provide information about the polarization. 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 the test set.



**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. We abbreviated the following glaciers: Columbia (COL), Dinsmoore, Bombardier, Edgeworth (DBE), Sjögren-Inlet (SI), Jakobshavn Isbrae (JAK).

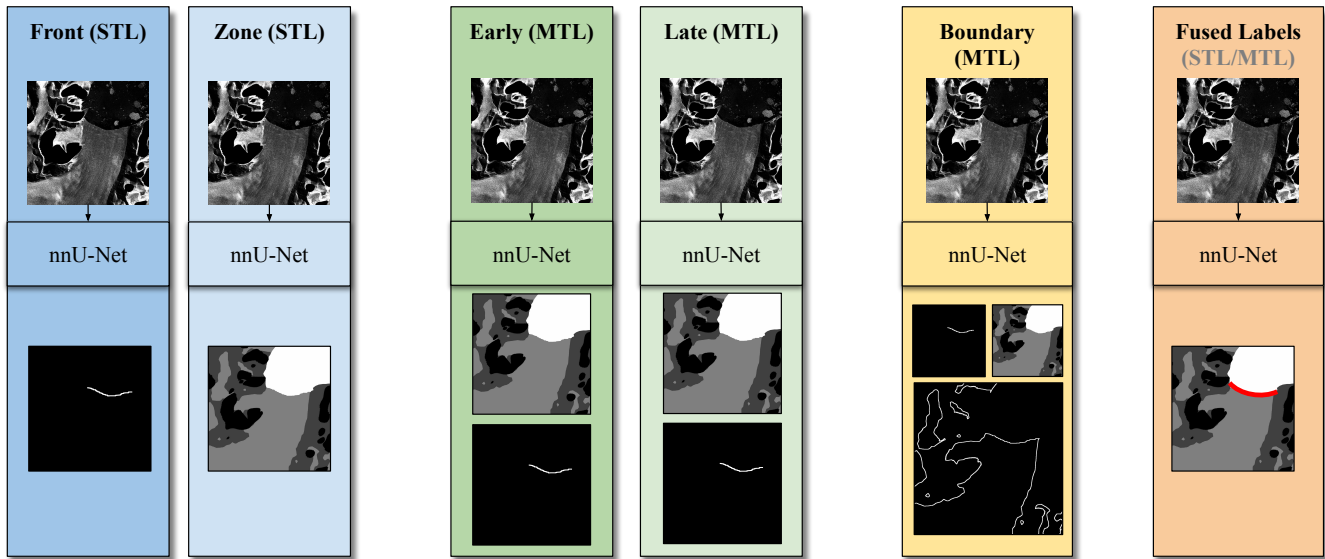
## 4 Method

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

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. The nnU-Net is a framework around the U-Net architecture. It provides good default values for hyperparameters, rules to adapt hyperparameters to the dataset, and many established deep-learning techniques for training and inference. Most of the rule-based hyperparameters are only relevant for three-dimensional (3D) data in the medical domain, like Computer Tomography (CT) and Magnetic Resonance Imaging (MRI) and are irrelevant to our experiments. ~~This dataset’s parameters that are important to its distribution and image quality determine its variability~~ ~~for all image modalities except CT, z-score normalization is applied. Each image is normalized independently~~ <sup>DI (1.2)</sup> by first subtracting its mean and then dividing by its standard deviation. Another dataset 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 and generalize well across various tasks. The nnU-Net uses a poly learning rate scheduler, a combination of dice coefficient and




**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).

160 cross-entropy as the loss function, and Stochastic gradient descent (SGD) with a Nesterov momentum as optimizer. The dice coefficient also helps with the problem of class imbalance. The coefficient of each class is weighted by the number of pixels in the label that relate to the class. 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. The foreground oversampling is especially helpful for our application because the class imbalance between the calving front and background is  
 165 high. It ensures that (at least) one-third of the patches for training are guaranteed to contain a foreground class. In our case, every batch contains two patches, from which one is forced to contain calving front pixels.

This framework achieves robust results for various medical image segmentation tasks. Overall, nnU-Net sets a new state of 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  
 170 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 segmentation, 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  
 175 blue columns on the left. The label of the calving front is dilated to the width of five pixels. Our preliminary experiments have



**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.

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.

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. 180 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 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 185 the framework. During the experiment’s planning, i. e., pre-processing phase, we fixed the network architecture’s estimated size to the size of the early-branching network so that late- and early-branching networks return the same value for the network size. <sup>A1 (1.2)</sup> Otherwise, early and late branching networks would be trained on different patch sizes. Thus, performance differences during the evaluation might arise from the different patch sizes and not the differing architectures. During training, the error of all labels is calculated and summed up with equal weighting. For the inference, we adapted only the number of channels in the case of 190 MTL. 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. We note that the prediction for respectively other task is omitted in the inference. <sup>A2 (1.5)</sup>

The last two experiments concern label changes. In Fig. 6, they are coloured yellow and orange. The fifth experiment of 195 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 third segmentation task is not considered in the final evaluation. The last experiment fuses the zone and 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 7x7 kernel.

### 4.3 Experimental Setup

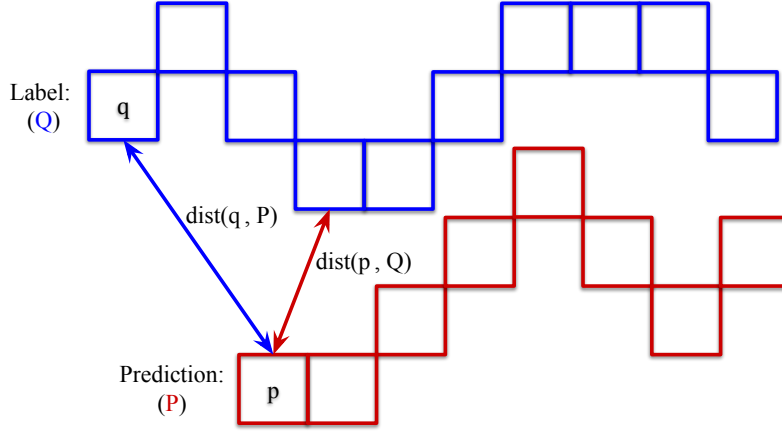
The nnU-Net was trained on a NVIDIA RTX 3080 with 12GB memory; the adapted network architecture has nine encoder blocks and eight decoder blocks. Each block consists of two convolutional layers: An instance normalization and a Rectified Linear Unit (ReLU). The kernels of all convolutional layers have a size of 3x3. During training, one batch contains two images. The patch size of the experiments that include only one label is 1280x1024. The experiments that have more than one label have a patch size of 1024x896 because of the GPU memory limit. There are also fixed parameters that are independent of the dataset. This includes the SGD optimizer with an initial learning rate of 0.01, a Nesterov momentum of 0.99, and a weight decay of 3e-5. Training of one epoch took between 100 s and 160 s. The maximum number of epochs of 500 is reached in every training (due to limited resources, we reduced the original maximum number of epochs from 1000 to 500). The nnU-Net defines one epoch using a fixed number of iterations (250). In each iteration, the batch is sampled depending on the class distribution of the sample to counteract the class imbalance.

## 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.

### 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 taken as the mean distance between the two lines (see Fig. 8 and Eq. 1). We note that the front MDE is also calculated for the



**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  $q$ .

zone segmentation. We extract the front from the zone segmentation by assigning all glacier pixels with neighbouring ocean pixels as the glacier front.

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

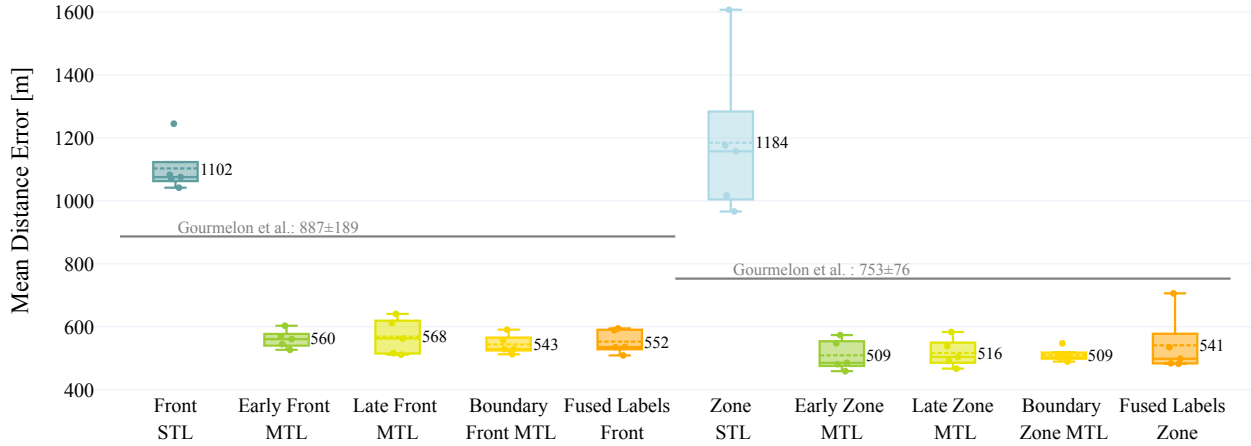
230 Additionally, we give classical segmentation metrics to evaluate the zone prediction. They include the Intersection over Union (IoU), which is defined as  $IoU = \frac{T_p}{T_p + F_p + F_N}$ , i. e., the True Positive ( $T_p$ ) pixels over the sum of  $T_p$ , False Positive ( $F_p$ ), and False Negative ( $F_N$ ) pixels. Additionally, the F1-score is computed ( $F1 = 2 \cdot \frac{pr \cdot re}{pr + re}$ ), which is a combination of Recall ( $re = \frac{T_p}{T_p + F_N}$ ) 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

235 In this section, we evaluate our experiments. In section 5.2.1, we compare the performance of the six experiments, which we described in section 4.2. In section 5.2.2, we analyse the prediction of the fused label experiment and investigate the influence of season, glacier site and satellite on the calving front prediction performance.

### 5.2.1 Comparison of the five experiments

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



**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.

samples that are incorrectly predicted to have no front pixel is similar, with an average of  $27.2 \pm 6.0$  for training on the front and  $24.8 \pm 12.4$  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,  $3.2 \pm 2.0$  for the front label out of 122 samples. The baseline of Gourmelon et al. is  $1 \pm 1$  based on the zone segmentation and  $7 \pm 3$  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  $887 \pm 189$  m for the front prediction and  $753 \pm 76$  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  $62.4 \pm 3.5$  IoU and  $71.7 \pm 3.2$  F1-score. The highest F1-score achieves the model that is trained with the additional boundary task with  $81.7 \pm 0.5$ . The model's IoU is  $72.6 \pm 0.4$ . The baseline of Gourmelon et al. is  $80.1 \pm 0.5$  F1-score and  $69.7 \pm 0.6$  IoU.


The fused label approach is the most feasible method discussed as follows. This approach performs on par with the other experiments. Even though the MDE of the boundary experiment is slightly lower than for the fused label approach, the student t-test showed that the difference is not significant. Moreover, the fused label approach requires fewer parameters and no changes



to the nnU-Net framework, making its use truly out-of-the-box. Additionally, this approach has a relatively small number  
 260 (5 ∈ 122) of predictions with falsely non-detected fronts.

### 5.2.2 Analysis of the fused label experiment

For the final evaluation, the zone segmentations of the five models that are trained during the five-fold cross-validation on the fused labels are combined into a more robust ensemble prediction. The models differ by their weight initialization and train-validation split. But are all trained on the fused labels. The ensemble prediction is created by taking the mean of the  
 265 probabilities for every class each model gives. We don't use the calving front pixels directly from the prediction but apply the post-processing of the zone label to gain a slightly better MDE. The post-processing labels the edge of the glacier zone as a calving front that has neighbouring ocean pixels. The performance of the ensemble prediction results in an MDE of  $515 \pm 39$  m and 5 ∈ 122 predictions with no fronts, which is similar to the average performance of the single models. We did not use the ensemble prediction for the comparison between the experiments in section 5.2.1 to show the variance that is introduced by a  
 270 different train-validation split and different weight initialization.



**Figure 10.** MDE on the test set grouped by glaciers and subdivided into seasons. The orange box plot in the first row shows all MDE of the test set. The olive green plot below shows the errors from the COL images, subdivided into seasons in the third and fourth rows. The last three rows show errors from the Mapple images. 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 ensemble model was trained with the fused zone and front labels. The y-axis has a logarithmic scale. The median is the middle line in the rectangle, and the dashed line represents the mean. The x-axis has a logarithmic scale. Otherwise, the outliers would dominate the plot. The rectangle reaches from the first quartile to the third quartile.


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 117 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  $287 \pm 48$  m for Mapple and  $840 \pm 48$  m for COL. The MDE of our methods is on average  $109 \pm 90$  m for Mapple, while the MDE of COL is  $930 \pm 1420$  m. The difference of the MDE is caused by a group of predictions with an error  $> 1000$  m. The median value is 222 m for COL and 80 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. 13).

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  $307 \pm 730$  m than the samples captured during the winter months ( $818 \pm 1389$  m). However, the medians are closer together with 150 m in the summer months and 181 m in the winter months. The MDE baseline (Gourmelon et al., 2022) in summer is  $732 \pm 93$  m and  $776 \pm 65$  m in winter. Although, most outliers are from winter seasons in COL. In winter, snow coverage of the glacier is more likely. The SAR signal can penetrate through several meters of snow cover, depending on the SAR frequency and the water content of the snow cover. However, most of the studied glaciers, in particular Columbia Glacier and the glaciers on the Antarctic Peninsula, are located in temperate marine environments making wet snow conditions also likely during winters. Thus snow can cover useful artefacts like crevasses and rock structures that can be useful for the pattern recognition of the nnU-Net. However, more important is the fact that the ocean next to the glaciers is covered more often by ice mélange during winters, which reduces the contrast between ocean and glacier areas. Even for experienced human mappers, it can be a challenging task to distinguish between the glacier tongue and the ice mélange. Thus, we hypothesize that the network has similar issues leading to reduced performance during winter. The difference between seasonal MDE also depends on the sensor. The MDE of the low-resolution sensors (S1, ENVISAT) is lower in the summer month. The MDE of high-resolution sensor (TDX) is lower in winter months but the difference is much smaller. A table of the MDE grouped by satellite and season is in Appendix A4.

An overview of the impact of the sensor resolution is given in Fig. 11. The images with a resolution of 7 m per pixel have a mean MDE of 180 m resolution have a mean MDE of 135 m, but this class contains only images of Mapple glacier taken by PALSAR. The images with 20 m per pixel have a mean MDE of 1214 m. The distribution has a cluster of MDEs that is larger than 1000 m. These large errors only stem from images of COL (see Fig. 12 row Columbia\_S1).

In Fig. 12, the MDE is grouped by satellites. The predictions for the samples captured by ERS, ENVISAT, and PALSAR have a similar average error between 150 m and 300 m. However, the samples from ERS, ENVISAT, and PALSAR only represent Mapple Glacier. TDX captures both test sites and has a MDE of  $180 \pm 179$  m. The samples captured by S1 have a higher MDE with  $1664 \pm 1807$  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 14 outliers have a front-line delineation error of  $> 1000$  m. These are predictions from images of COL taken by S1. The low-resolution, e. g., cracks in the ice, are not visible. Fig. 13 (c) shows that the model falsely predicts ocean pixels as the front and does not predict the left calving front at all. The MDE for this sample is  $4221 \pm 5026$  m. We suspect that the complex calving front of COL requires a high resolution for accurate detection or more training data of S1. The group of outliers heavily increases the overall MDE but can be separated clearly by specifying satellite and glacier. It shows

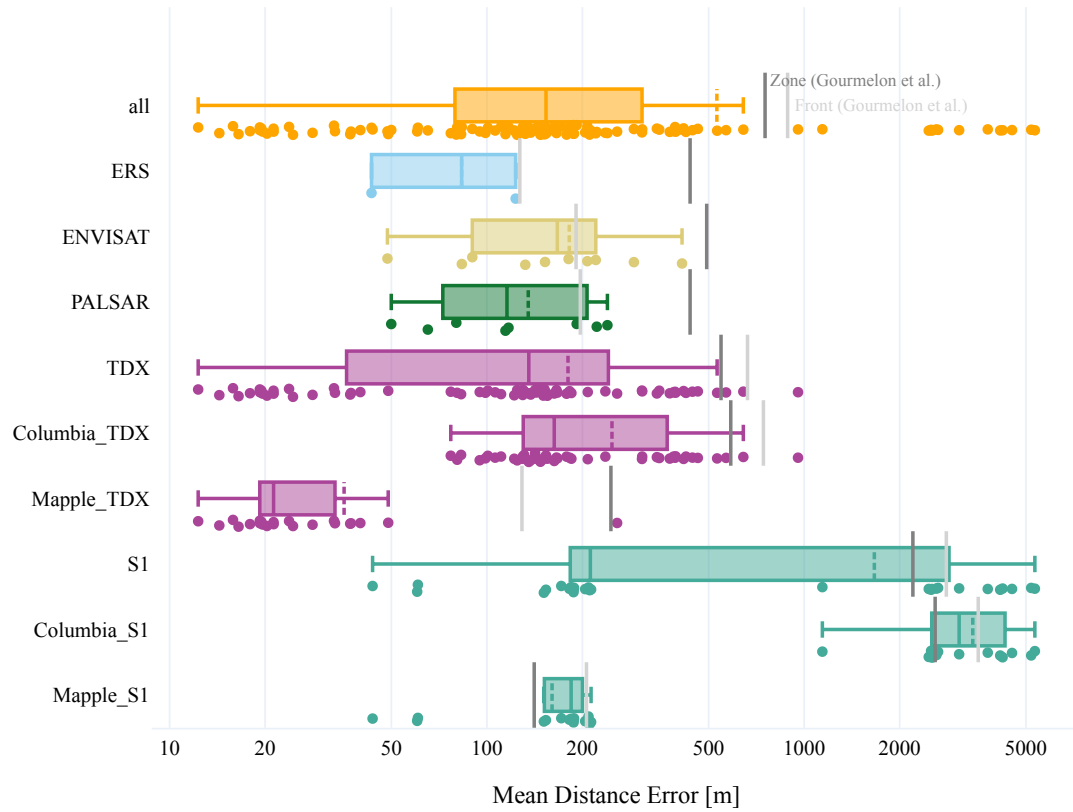


**Figure 11.** MDE on the test set grouped by resolution of the SAR image. The orange box plot in the first row shows all MDE of the test set. ENVISAT, ERS and S1 have a resolution of 20 m per pixel. MDE of this resolution are represented by the dark red box plot. PALSAR has a 17 m per pixel resolution. MDE of images of this resolution are represented in the magenta box plot. TDX has a 7 m per pixel resolution. MDE of images of this resolution are represented in the green box plot. 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. The median is the middle line in the rectangle, and the dashed line represents the mean. The x-axis has a logarithmic scale. Otherwise, the outliers would dominate the plot. The rectangle reaches from the first quartile to the third quartile.


the generalisation limits of our method. For a visual inspection, we provide the predictions of the test set as single images and as an animation on Zenodo (<https://zenodo.org/record/7915740#.ZFtS3dlzZhE>) (<https://zenodo.org/record/8379954>).

310 The MDE is negatively influenced by calving front predictions extending over the labelled calving front on the coastline, which can be seen in Fig 13 (a) and (b). The MDE might give a wrong implication on the usability of these predictions. The prediction of 13 (b) has a relatively high MDE (393 m) because it predicts the coastline as a calving front. Despite this, the prediction can still be used to monitor the change in the position of the calving front as well as the calculation of the total frontal ablation rate or surface area. The prediction can be corrected easily if the coastline is known. The misclassification of the coastline as a calving front holds true for nearly all images of COL taken by TDX. The northern part of COL is particularly difficult to distinguish between the coastline and calving front because the glacier retreated so far that it transitioned from a marine-terminating glacier to a land-terminating glacier. Without prior knowledge, it is even difficult for humans to distinguish between the coastline and the calving front.


320 The model predictions of the calving front on images of Mapple have a high overlap with the label, even with a low resolution of  $\geq 20$  m (ERS) (see Fig. 14). 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. 14).



**Figure 12.** MDE on the test set grouped by satellites. 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 images from ERS (light blue), ENVISAT (yellow), and PALSAR (green) only capture the Mapple Glacier. The MDE from TDX (pink) and S1 (turquoise) is subdivided into the two glaciers of the test set. The ensemble model was trained with the fusion of zone and front labels. The y-axis has a logarithmic scale. The median is the middle line in the rectangle, and the dashed line represents the mean. The x-axis has a logarithmic scale. Otherwise, the outliers would dominate the plot. The rectangle reaches from the first quartile to the third quartile.



**Figure 13.** Calving front predictions of COL on (a) 11 February 2016, and (b) 22 June 2014 taken by TDX with 7m pixel resolution (after multi-looking with a factor of 3x3); (c) 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 14.** 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 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 ENVISAT on 22 September 2007. SAR imagery is provided by DLR, ESA, and ASF.

## 6 Conclusion

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 increases each task's performance. No significant difference between the two Multi-Task-Learning architectures exists. 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  $541 \pm 84$  m. We provide the code and the pre-trained model for application on further SAR images of glacier fronts (see Code and data availability). The predictions need to be filtered manually since there can be outliers, as our results show. But it will provide an initial prediction, which eases the task of glacier front delineation.

To improve the average MDE, future work should focus on reducing the model performance on low-resolution images like S1 with 20 m per pixel resolution. This can be done by including more images taken by S1 in the training set or by implementing an oversampling strategy of the low-resolution images.

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 experiments, where only one label is used. 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. Our findings highlight the suitability of the nnU-net for glacier front segmentation on multi-mission SAR remote sensing data, which will facilitate an efficient, extended spatiotemporal mapping of tidewater glacier terminus changes. Our findings also promote the out-of-the-box application of the nnU-Net for other segmentation tasks based on satellite imagery, because we didn't need to modify it for the calving front detection.

## 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 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. We provide Fig. table A5 to compare the two methods (nnU-Net and (Gourmelon et al., 2022)). Table A4 contains the MDE of the fuse label training grouped by satellite and season.

	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
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.

## 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 colored with a gradient from bright for the past and dark red for the most recent fronts.

## Appendix C: Temporal distribution of zone label

355 Figure C1 show the temporal distribution of the zone label. 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



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

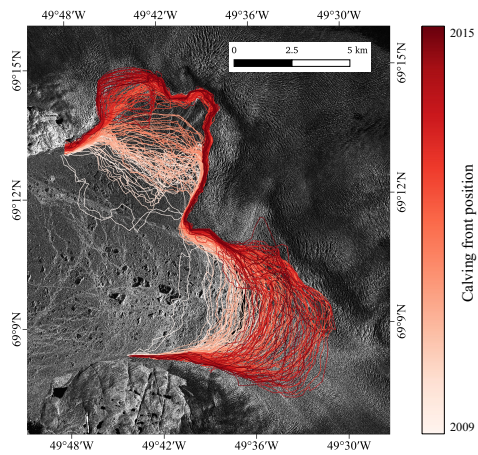
Satellite:	S1	ENVISAT	ERS	PALSAR	TDX	all
winter	2529 $\pm$ 1719 m	241 $\pm$ 101 m	-	-	160 $\pm$ 125 m	818 $\pm$ 1389 m
summer	798 $\pm$ 1442 m	122 $\pm$ 61 m	83 $\pm$ 39 m	135 $\pm$ 68 m	197 $\pm$ 213 m	307 $\pm$ 730 m

**Table A4.** MDE of the fused label experiments grouped by sensor and season.


information available class. JAK’s area distribution shows a repetitive structure of increasing and decreasing glacier area. This pattern represents seasonal changes.

Hyperparameter	Gourmelon et al. (2022)	nnU-Net (Isensee et al., 2021)
Number of convolutional layers	10	34
Pooling	Max-pooling	Strided convolution
Activation-function	ReLU	Leaky ReLU
Patch size	256x256	1280x1024
Batch size	16	2
Optimizer	cyclic learning rate scheduler (Smith, 2017) in combination with the Adam optimizer (Kingma and Ba, 2014)	SGD with a Nesterov momentum of 0.99, and a weight decay of 3e-5
Initial learning-rate	$4 \cdot 10^{-5}$	$1 \cdot 10^{-3}$
Gradient clipping	True	False
Deep-supervision	False	True
Loss-function	Combination of Dice and Cross-entropy	Combination of Dice and Cross-entropy


**Table A5.** Comparison of the U-Net training setup and hyperparameters between Gourmelon et al. (2022) and the nnU-Net (Isensee et al., 2021).




**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.



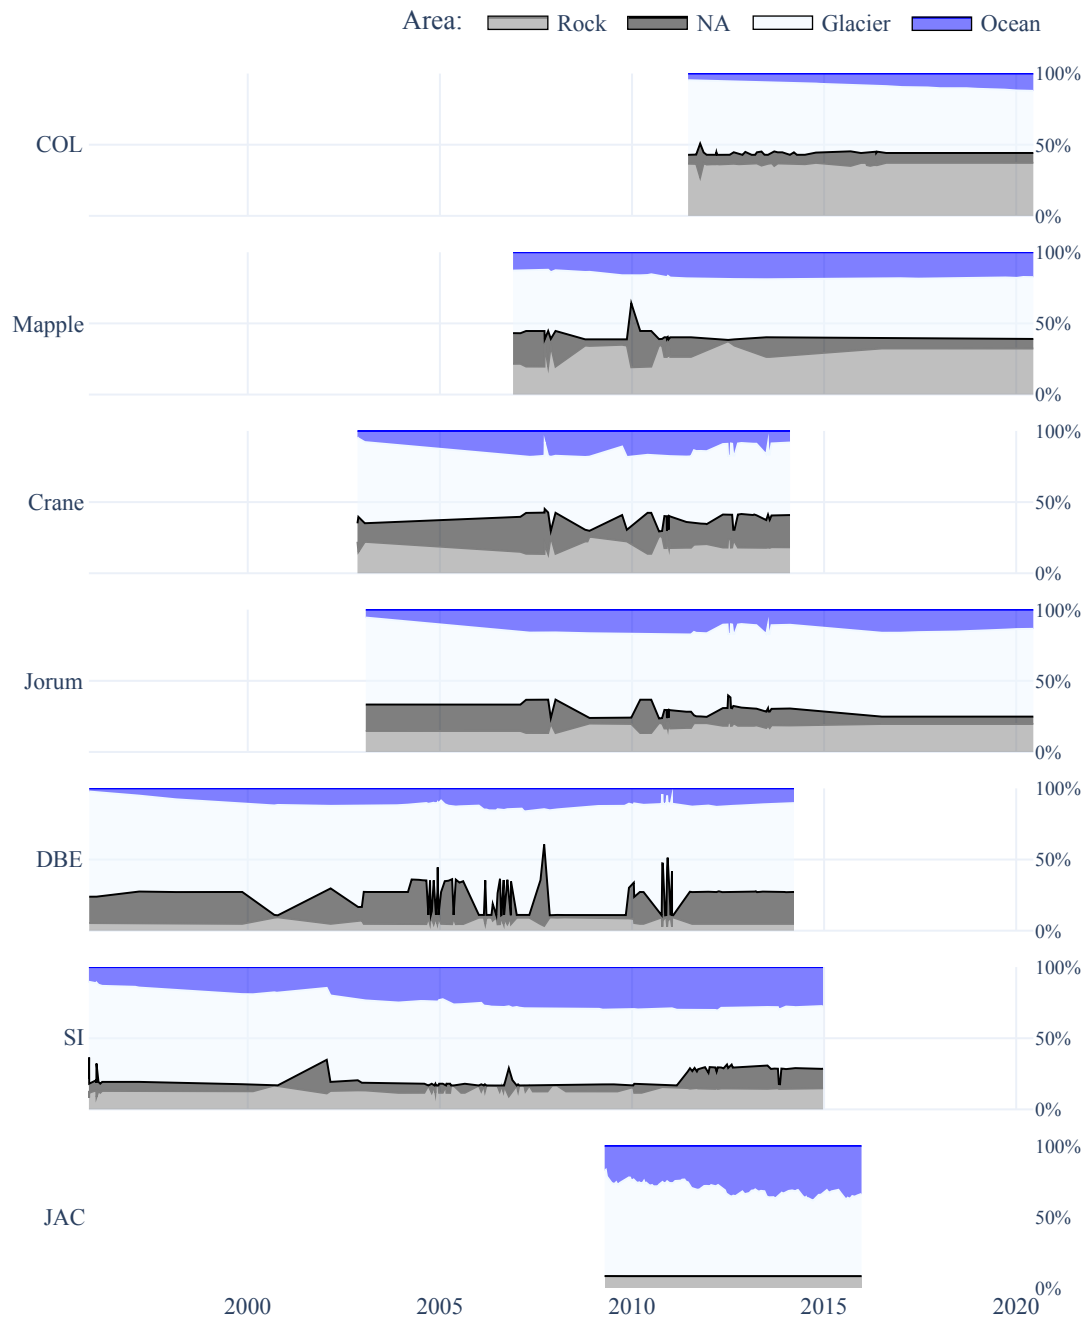
**Figure B4.** DBE Glacier from 1995 to 2014.



**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.

## 360 **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 visualization scripts is available on GitHub

([https://github.com/hol11laqe/nnUNet\\_calvingfront\\_detection](https://github.com/hol11laqe/nnUNet_calvingfront_detection)) and a Demo-version is provided on HuggingFace

365 ([https://huggingface.co/spaces/hol11laqe/nnUNet\\_calvingfront\\_detection](https://huggingface.co/spaces/hol11laqe/nnUNet_calvingfront_detection)). The calving front predictions of the test set are available on Zenodo (<https://zenodo.org/record/7915740#.ZFtS3dIzZhE>) (<https://zenodo.org/record/8379954>) as well as the <sup>C3 (1.1)</sup> pre-trained model

(<https://zenodo.org/record/7837300#.ZFuQAdIzbUA>). The dataset is provided by Gourmelon et al. (2022)

(<https://doi.org/10.1594/PANGAEA.940950>).

## 370 **Author contribution**

OH and NG were responsible for conceptualizing the study. OH developed and implemented 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. JF acquired financial support. OH created visualizations and prepared the manuscript with contributions from all co-authors. JF, VC, and NG reviewed and edited the manuscript.

## 375 **Competing interests**

At least one of the (co-)authors is a member of the editorial board of The Cryosphere.

## **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 by the ESA

380 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 programme "Open Access

385 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 "Masuring and modelling mountain glaciers and ice caps in a changing climate (M<sup>3</sup>OCCA)".

## References

- 390 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>, 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, 104037, <https://doi.org/10.1016/j.combiomed.2020.104037>, 2020.
- 395 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.
- 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.
- 400 Baumhoer, C. A., Dietz, A. J., Kneisel, C., Paeth, H., and Kuenzer, C.: Environmental drivers of circum-Antarctic glacier and ice shelf front retreat over the last two decades, *The Cryosphere*, 15, 2357–2381, <https://doi.org/10.5194/tc-15-2357-2021>, 2021.
- Baumhoer, C. A., Dietz, A. J., Heidler, K., and Kuenzer, C.: IceLines – A new data set of Antarctic ice shelf front positions, *Scientific Data*, 10, 138, <https://doi.org/10.1038/s41597-023-02045-x>, 2023.
- 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.
- 405 Bischke, B., Helber, P., Folz, J., Borth, D., and Dengel, A.: Multi-Task Learning for Segmentation of Building Footprints 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.
- 410 Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H.: Encoder-Decoder with Atrous Separable Convolution for Semantic Image 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](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, Springer International Publishing, Cham, [https://doi.org/10.1007/978-3-030-32248-9\\_51](https://doi.org/10.1007/978-3-030-32248-9_51), 2019.
- 415 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.
- 420 Cook, A. J., Murray, T., Luckman, A., Vaughan, D. G., and Barrand, N. E.: A new 100-m Digital Elevation Model of the Antarctic Peninsula derived from ASTER Global DEM: methods and accuracy assessment, *Earth System Science Data*, 4, 129–142, <https://doi.org/10.5194/essd-4-129-2012>, 2012.
- 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>, 2022a.
- 425

- Davari, A., Islam, S., Seehaus, T., Hartmann, A., Braun, M., Maier, A., and Christlein, V.: On Mathews Correlation Coefficient and Improved Distance Map Loss for Automatic Glacier Calving Front Segmentation in SAR Imagery, *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–12, <https://doi.org/10.1109/TGRS.2021.3115883>, 2022b.
- 430 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.
- 435 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 440 Relevance in Multi-Task U-Net for Accurate Prostate Segmentation in CT Images, *IEEE transactions on medical imaging*, 40, 2118–2128, <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., 445 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.
- 450 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- 455 aware multi-label chest X-ray organ segmentation, *International Journal of Computer Assisted Radiology and Surgery*, 15, 425–436, <https://doi.org/10.1007/s11548-019-02115-9>, 2020.
- Kingma, D. P. and Ba, J.: Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980*, 2014.
- 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.
- 460 Li, X., Wang, Y., Tang, Q., Fan, Z., and Yu, J.: Dual U-Net for the Segmentation of Overlapping Glioma Nuclei, *IEEE Access*, 7, 84 040–84 052, <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.
- 465 Marochov, M., Stokes, C. R., and Carbonneau, P. E.: Image Classification of Marine-Terminating Outlet Glaciers using Deep Learning Methods, preprint, *Glaciers/Remote Sensing*, <https://doi.org/10.5194/tc-2020-310>, 2020.
- 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.
- Minowa, M., Schaefer, M., Sugiyama, S., Sakakibara, D., and Skvarca, P.: Frontal ablation and mass loss of the Patagonian icefields, *Earth and*  
 470 *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.
- 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,  
 475 <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.
- 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.
- 480 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](https://doi.org/10.1007/978-3-319-24574-4_28), 2015.
- 485 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, 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,  
 490 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., Mouginit, 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, 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,  
 495 <https://doi.org/10.1038/s41586-018-0179-y>, 2018.
- Smith, L. N.: Cyclical Learning Rates for Training Neural Networks, in: *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 464–472, <https://doi.org/10.1109/WACV.2017.58>, 2017.
- 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



- 500 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 fronts of  
505 Greenland glaciers from multi-sensor remote sensing imagery, *Remote Sensing of Environment*, 254, 112 265, <https://doi.org/10.1016/j.rse.2020.112265>, 2021.