Analysis of Ice Shelf Front Dynamics in Pine Island Bay (Antarctica) based on long-term SAR Time Series and Deep Learning

LUISA WAGNER[123], CELIA BAUMHOER[2] & TOBIAS ULLMANN[1]

Abstract: Ice shelves, the floating extensions of ice sheets, create a safety band around Antarctica by buttressing the upstream grounded ice. Loss of their stability leads to increased ice discharge contributing to sea level rise. Thus, it is important to monitor ice shelf dynamics. So far, the potential of spaceborne Synthetic Aperture Radar (SAR) data has not yet been full exhausted as early satellites have hardly been used. To fill this research gap, this study made use of the ERS and Envisat archive within Pine Island Bay, a region of drastic ongoing changes. A 20-year time series (1992-2011) of ice shelf front dynamics was derived using a deep neural network architecture that combines segmentation and edge detection. The product reveals individual destabilisation patterns - most considerable for Pine Island Ice Shelf in terms of frequency of calving events and Thwaites ice tongue in terms of size of break-up. The successful approach demonstrates the potential for long-term analysis of ice shelf dynamics over entire Antarctica.

1 Introduction

The Antarctic Ice Sheet stores a volume of water equivalent to a potential global sea level rise by 58 m. Although so far, the contribution is much smaller (ca. 0.3 mm/year), the current acceleration of mass losses indicates that impacts to humankind will increasingly become threatening on global scales (IMBIE 2018). Ultimately inevitable adaptation measures demand precise predictions of the rate of sea level rise. This requires knowledge of mechanisms affecting the mass loss, which is controlled by ice shelves, the floating extensions of ice sheets. Ice shelves almost equally lose mass by basal melt and iceberg calving changing the ice shelf front position. This has fundamental consequences for the upstream ice sheets, because ice shelves restrain the ice flow by resistance from lateral forces through confining margins and local points of grounding. A reduction of this stabilising buttressing effect following a calving event leads to acceleration of ice flow and increased ice discharge into the ocean (FÜRST et al. 2016).

A key parameter for process understanding and improved predictions is therefore the position of the calving front that can be monitored with remote sensing data. But the tedious manual front delineation and data limitations did not allow large scale inter-annual calving front analysis in the past. Today, SAR sensors provide plenty of earth observation data even during polar night and the

[1] University of Würzburg, Department of Remote Sensing, Würzburg, Germany
[2] German Aerospace Center, Earth Observation Center, Oberpfaffenhofen, Germany
[3] Bavarian Academy of Sciences and Humanities, Geodesy and Glaciology, Munich, Germany, E-Mail: luisa.wagner@badw.de

presence of clouds. But extracting valuable information from complex SAR data over the Antarctic coastline requires new image processing techniques such as deep learning-based approaches. This complexity exists because of the challenging separability between the ice shelf and its surroundings – especially the sea ice (BAUMHOER et al. 2018). These issues have been successfully tackled with deep learning applications on Sentinel-1 data (BAUMHOER et al. 2019). However, early SAR sensors have up to now only been used to a limited extent. Only few studies use single scenes, but none exploit the full potential of available archive data. Given that these sensors deliver information about much longer time scales, this is a huge research gap. This study addresses this gap by taking advantage of ERS and Envisat SAR data for the purpose of automatically delineating ice shelf fronts in Western Antarctic Pine Island Bay. The research questions target (1) the suitability of the data, (2) the transferability and adaptation of existing deep learning frameworks to new data and (3) the potential of derived time series to reveal ice shelf front dynamics.

2 Study Area

The Western Antarctic Pine Island Bay (75°S 102°W) is located in the Amundsen Sea, an arm of the Southern Ocean. The coast is extensively fringed by ice shelves, namely Cosgrove (CGIS), Pine Island (PIIS), Thwaites (TIS), Crossen (CIS) and Dotson Ice Shelf (DIS). They constitute a major ice drainage basin of the Western Antarctic Ice Sheet. The landward-sloping topography beneath the grounded ice results in an inherent instability that suggests further acceleration of already observed mass loss trends connected to fast ice flow, ice shelf thinning and grounding line retreat (MACGREGOR et al. 2012). While these changes are considerable for all investigated ice shelves relative to their respective size, they are exceptionally high for PIIS and TIS (LHERMITTE et al. 2020). This is especially concerning as they are already the biggest contributors to current global sea level rise (~5%). Although ocean warming, resulting in enhanced basal melt rates and ice shelf retreat, has been identified as major responsible process, large modelling uncertainties concerning their future dynamics exist (ALLEY et al. 2021).

3 Data and Methods

3.1 Data and pre-processing

The study is based on C-band SAR data from the satellites ERS-1, ERS-2 (both HH polarisation), Envisat (HH and VV) and Sentinel-1 (HH and HV). All datasets are used in the model training, while only ERS and Envisat data are subject to inference. The Sentinel-1 dataset is taken from a study by HEIDLER et al. (2021). For ERS and Envisat, all scenes available in the archive (except duplicates) are downloaded. This provides a data base of 1990 scenes with a temporal coverage from 1992 to 2011. Pre-processing includes radiometric calibration, speckle-filtering with a refined Lee-filter and illumination terrain correction with the Copernicus DEM. For the sake of comparability and data reduction, all scenes are bilinearly interpolated to a spatial resolution of 30 m and cropped to a buffer of 50 km around an approx. coastline from 1997. Three different enhancement techniques were tested to identify a best practice approach. The comparison included no enhancement, percentile clip followed by normalization or batch normalization. While the

second removes outliers and rescales backscatter values to a common range, the latter removes internal changes in the distribution of model network parameter values. The prepared scenes are tiled into patches (768 × 768 pixels). To avoid noise at the edges and to increase the amount and variability, neighbouring tiles overlap by 50% and eightfold augmentation (mirroring, rotating) is applied. The procedure results in a total amount of 14,626 tiles.

3.2 Methods

A slightly modified version of the HED-UNet proposed by HEIDLER et al. (2021) is used in this study. It combines the tasks of segmentation and edge detection matching the intuitive way how a human annotator would approach the task of ice shelf front delineation. To learn both tasks synergetically, the model applies a multi-task learning technique that combines the essential advantages of a segmentation building block and an edge detection building block (Fig. 1). The segmentation building block is based on an encoder-decoder architecture, which is inherited from the underlying UNet model developed by RONNEBERGER et al. (2015). The encoder consists of a series of downsampling steps that aggregate contextual information at a lower resolution. This is followed by upsampling steps in the decoder that bring back the information to the pixel level. To enable the localisation of the pixel-wise prediction, high resolution features from the encoder are combined with their counterpart in each upsampling step. This procedure preserves the contextual information and is thus the architecture's main advantage. In comparison to the original UNet architecture, the number of used resolution levels is increased to six in order to enhance the spatial context.

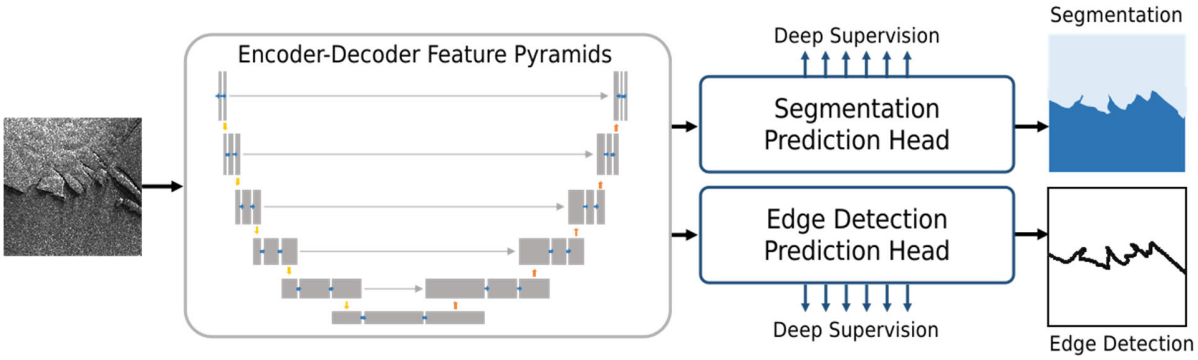

Fig. 1: HED-UNet architecture (modified after HEIDLER et al. 2021)

The edge detection building block based on the HED network (XIE et al. 2015) consists of an encoder that can be implemented in the same way like in the segmentation building block. Instead of stepwise decoding through each resolution level to the initial input resolution, the HED network uses a merging block that combines information from all resolution levels in the encoder. At each resolution level, a feature map is created by convolution and interpolated to the original image size. The resulting six images contain slightly different information that yield a combined prediction from all resolution levels. The high-resolution levels provide crisper, but less robust predictions. The lower resolution predictions are more robust, but at the expense of precision. Thus, the merging head combines sharp edges at the front with coarse information from the border areas. That avoids false positives, which constitutes the main strength of this edge detection block.

The multi-task approach combines the advantages of both tasks' blocks in a joint architecture. This is implemented through an encoder-decoder structure, where the encoder serves the two tasks simultaneously. This is followed by two task specific prediction heads in which the idea of multi-level prediction from the edge detection block is picked up. It is further enhanced by a hierarchical attention mechanism that enables that different fusion mechanisms can come into effect in different parts of the images. Near the coastlines, the predictions require information from the high-resolution layers, while farther away more general, coarser information are sufficient. The proceeding "allows the model to attend to different resolution levels" (HEIDLER et al. 2021) in the manner that it can focus on those features that are considered most important for each pixel's prediction. This is realised by respective weighting of the feature maps as opposed to using fixed weights for predictions. The two prediction heads are collaterally implemented for both tasks, but follow the same structure and independently generate the particular output they are targeting.
In order to train the model, inputs with known outputs are presented to the model during training. The performance is evaluated by a class-weighted cross-entropy loss function and network parameters are iteratively adjusted to get improved outputs. Following the principle of multi-resolution prediction, the learning effectiveness is enhanced by adding additional intermediate supervision at different resolution levels of the encoder-decoder structure as opposed to only using one loss term for the entire network. Downsampled versions of the ground truth tiles are compared with the respective layer's intermediate prediction. This way, more direct feedback is provided in the form of additional loss terms that support learning of progressively more suitable weights in the next layer. The training is based on manual labels of 27 temporally and spatially diverse ERS and Envisat scenes and further enhanced by existing Sentinel-1 training data. The total number of 92 labels is divided into a training and a validation subset. Two different training approaches are tested – from scratch and on-top of an existing model (training is initialised with existing weights instead of using random initial weights). This can compensate for deficient availability of training data. The used pre-trained weights origin from a model by BAUMHOER et al. (2023) that is based on Sentinel-1 data between 2015 and 2022 located all around Antarctica. Due to different polarisations, the pre-trained model has to be slightly adapted in order to train on-top with single polarised ERS and Envisat scenes. The training for all experimental set-ups was conducted for 30 epochs on a Global Processing Unit (GPU) of type "NVIDIA GeForce GTX TITAN X" with a batch size of four and a training rate of 0,0001.
The post-processing workflow mosaics the inferred results to obtain regular results despite uneven data availability. Depending on data availability, yearly, seasonal and monthly products are created by averaging the prediction probabilities of all scenes. A morphological filter removes holes in the form of small patches of misclassified pixels. DEM based correction is applied to automatically classify all pixels high above ice shelves as ice to remove further potential misclassified pixels. The results are converted to line features to obtain the final calving front.
The accuracy of the predictions is assessed using both a segmentation and a distance measured accuracy assessment scheme. They are based on selected manually depicted reference fronts from different locations and timings. To avoid bias introduced by robust predictions far away from the front, the accuracies are calculated in a 3 km buffer around the manual reference front.

4 Results

4.1 Accuracy assessment

Accuracy comparison of all tested image enhancement techniques revealed that batch normalisation brings enhancement whereas percentile clipping plus normalisation does not. Different training set-ups demonstrated that training with Sentinel-1 data in addition to ERS and Envisat cannot improve the results. The choice of training approach, however, has a huge impact. On-top training improves the accuracies by 10%. Consequently, the on-top trained model without enhancements was identified as performing best with all accuracy measures exceeding 90%. The average accuracy yields 92.2% and the f1-score 92.0%. The precision is slightly higher (93.6%) than the recall (90.8%). The distance measured accuracies yield a mean deviation of 354.95m.

4.2 Ice Shelf Front Dynamics

The extracted front positions reveal individual fluctuation for all investigated ice shelves (Fig. 2). Given their extraordinary importance, the following description and analysis focuses exemplarily on the dynamics of PIIS and TIS. The other ice shelves exhibit only very little overall retreat or advance, only CIS is characterised by intermediate unsteady variations.

PIIS, although not showing the fastest advance rates in Pine Island Bay, was found to be the most dynamic in terms of calving frequency. Multiple calving events increasing in retreat rates followed by re-advance occurred at regular intervals (Fig. 2, close-up C). The first small one in 1995 accounts for 3.3 km of retreat (along the centreline), the next in 2007 for more than 10 km and in the last in 2007 for more than 15 km. The interim rates of advance also increased from 1707 m to 3210 m per year. Thus, the overall front position did not change much from the beginning to the end of the study period, but the absolute total retreat is considerable with more than 30 km. Another development that can be observed in the satellite imagery time series is fracturing at the margins, leading to more pronounced indentations.

The investigation of TIS has to focus on two parts. The ice tongue was found to exhibit the largest dynamics in the study area in terms of the magnitude of change, while probably the lowest concerning the frequency. Just one very huge calving event occurred in 2002 which led to a retreat by more than 80 km. The following re-advance (Fig. 2, close-up A) slowed down to ~3500 m per year compared to an average advance of ~6300 m per year before the break-up. The re-advance was very static and the shape of the break-off edge continued to define the shape of the new ice shelf for years, as no further calving occurred. The western part of the front that was positioned slightly further back at the beginning, caught up with the northern side throughout the years. The lateral boundary of the ice shelf is very rough. On the side bordering the Eastern Ice Shelf, an indentation widened and a second small ice shelf branch starts to advance in 2010 and 2011.

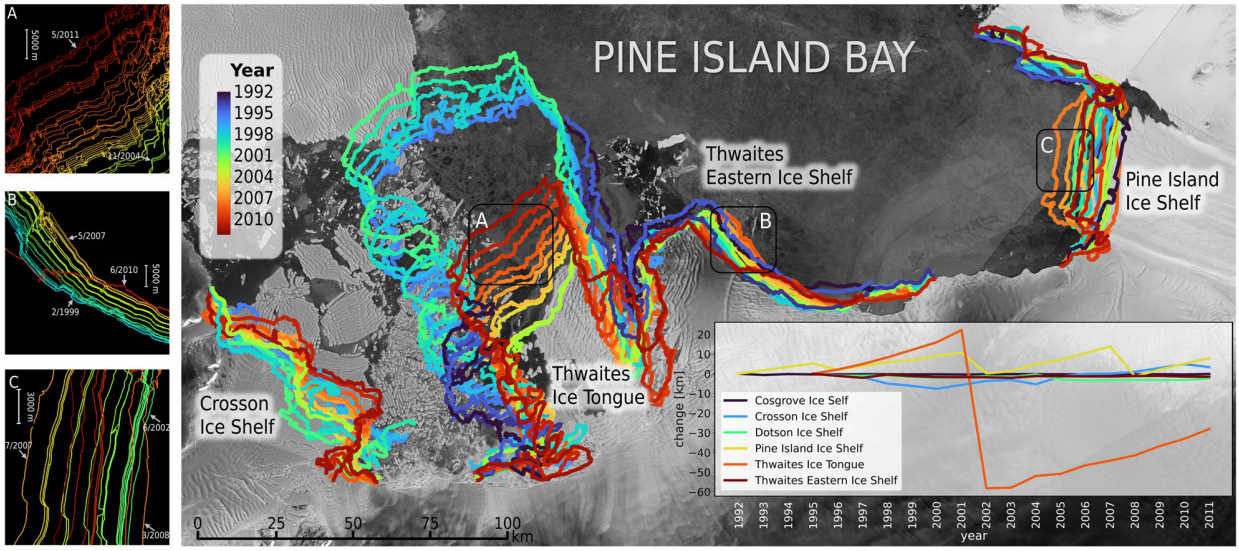

Fig. 2: Map of yearly calving front positions of selected ice shelves. Close-up maps A, B and C depict monthly positions of selected areas. The plot in the lower right shows the time series of calving front change along the centrelines for all investigated ice shelves

The front of the Eastern Ice Shelf showed much less dynamics. It is slightly fluctuating in the 1990s, but remained mainly stable in the 2000s. Changes at the lateral edges are much more pronounced. This refers to the indentation at its western part connecting the ice shelf to the ice tongue which progressively widens inward. Simultaneously, on the eastern side, a steady advance started in the mid 1990s (Fig. 2 close-up B). A lateral calving event resulted in a retreat by 6.8 km. At the same time the westward indentation spread further into the ice shelf.

5 Discussion

5.1 Accuracy assessment

The identified best model stands out concerning the overall performance (measured through the f-1 score) and the precision, i.e. the predicted positives are mostly correct. A detailed accuracy assessment of this model confirms this finding. However, it becomes apparent that considerable differences between the individual testing scenes exist depending on the complexity of the scene. Roughed and mélange prone ice shelves exhibit generally lower accuracies (e.g. CIS) and accordingly dominate the evaluation. Nevertheless, the on average measured uncertainties are much smaller than the detected dynamics and the performance lies within a comparable range like other calving front studies (e.g. PERIYASAMY et al. 2022). Superiority of the applied method to formerly common approaches seem to balance the disadvantages of older SAR data. These disadvantages include cases of geolocation errors in the raw data and seasonal data irregularities or even complete data gaps.

The developed post-processing workflow can, to some extent, counteract those issues by removing artefacts, e.g. at the border of the scene or in the form of holes within a predicted class. The approach of mosaicking as part of the post-processing also has a positive effect on coping with

seasonal backscatter variations. This is particularly important in summer months when enhanced melt can lead to significant changes of the signal with potential reduced separability of ice shelves and ocean, or sea ice, respectively. However, the mosaicking also has side-effects that have to be considered during interpretation of the results. Issues may occur in the case of huge dynamics within the composited time. The mosaic combines the information about the front position before and after a calving event and the merged probabilities between the two fronts are reduced compared to the stable ice shelf parts. Accordingly, the accuracy of the automatic binarisation is limited. In order to deal with this, additional seasonal and monthly composites have to be obtained. This is also advantageous in light of some more general considerations of the approach. On the one hand, mosaicking simply reduces the frequency of results. On the other, it can introduce ambiguities by aggregating over a time interval. The actually represented timing of results based on mosaics is not clear. To somewhat assist interpretation despite this issue, the calculation the mean day of the year of each pixel in a mosaic can be consulted. This way the yearly results are sufficient to detect general trends and the occurrence of calving events. Seasonal and monthly composites minimise the problem, but their availability is irregular, i.e. the non-equidistant time steps have to be considered. This shows that the main issue and challenge is the sparse and irregular data availability.

5.2 Ice Shelf Front Dynamics

Despite the discussed issues, the resulting time series prove to deliver valuable insights into ice shelf dynamics. A close look at the changes provides indications of developments and underlying processes. At PIIS, the ice shelf front shape after calving events indicate the break-up of tabular icebergs with lengths of several kilometers and widths close to the width of the ice shelf. The periods of advance are of comparable length, and it can be noted that the advance rate is faster, the bigger the previous calving event. The time series of the calving cycles include hints of topographic influences on changes, especially through investigation of the constantly changing fractured shear margins. The observed features are very likely rifts that develop due to high strain rates (RIGNOT 2002). These come into effect because surface velocities are much higher in the central ice shelf than in the marginal parts, which results in shear strain along the side and causes lateral drag (HAN et al. 2016). The rifting starts with the calving event in 1996 and continuously shifts and spreads into the interior of the ice shelf. The variations might be connected to ice flow around an ice rise and subsequent unpinning and ungrounding of the shelf (BINDSCHADLER 2002). The emerging fracture zones and deepening indentations constitute weak areas that play a major role in the later calving events. During the end of the observation period the edges of the ice shelf are increasingly disconnected from the coastline. This reveals parallels with the ice shelf geometry prior to previous calving events, which indicates continuation of the cycle of iterating advances and calving. Further, the rifting induced pattern suggest reduced side-shear resistance to ice shelf flow and inland ice flow, i.e. a reduction of the buttressing effect (RIGNOT 2002). The results thus do not only indicate frontal fluctuations, but also destabilisation of the ice shelf.

At TIS, the most obvious event within the study period is the calving of a huge tabular iceberg at the ice tongue in 2002. An exact dating is insufficient based on the available data, but a timing early in the year can be assumed. With a retreat of more than 80 km it is by far the hugest break-

off detected in the study area. Investigation of the geometry of dynamics, however, hints that this event is not the only drastic process. It rather can be interpreted as one component in a broader process of destabilisation. In the early years, the ice tongue exhibited the largest expansion and slightly less roughed margins than at any later point in the study period. It was described as structurally intact due to being grounded at a subsea ridge (WILD et al. 2022; TINTO et al. 2011).
Small details in the time series following the break-up hint that this changed at some point before the event. The assumption arises from the observation that the break-off edge is not close to the grounding line and no acceleration of the ice stream was detectable after the event. This manifests in the time series in the form of reduced advancement rates following the event. If pinned to the ground, an ice stream acceleration would be expected due to the reduced buttressing effect. Since this is not observed, it is an indication that the calved part was already unpinned and passive ice, i.e. already did not contribute to buttressing at that time (FÜRST et al. 2016). Investigations of the SAR scenes and the geometry of the front show that the new advancement is characterised by very fractured conditions on the remaining ice tongue. Comparison to the positions of known pinning points confirms that rifting and fracturing is most pronounced there. Multiple studies report that the contact to the pinning points becomes ephemeral in the 2000s (TINTO et al. 2011). Yet, the advance happened steadily and without major calving. From that can be drawn that weakening of the ice tongue did not primarily take place near the front, but mainly occurred in the part closer to the grounding line. It is especially pronounced in the shear zone to the bordering eastern ice shelf (LHERMITTE et al. 2020). A pronounced indentation during the last two years is spreading inwards up to a position much closer to the grounding line. On the western side of the ice tongue, it is even visible that the spreading is no longer only appearing opposite to the flow direction, but also starts to move perpendicular to it. The propagation seems to increasingly cut off the tongue. According to MACGREGOR et al. (2012), this is a sign of a calving event that took place without being very visually striking because the calved part is still connected by mélange. Thus overall, the detected dynamics of Thwaites ice tongue during the study period are indicative of a transition from a mainly intact state to a weaker state manifested by mélange of fractured icebergs (MILES et al. 2020).
The development of the eastern ice shelf is connected to the ice tongue weakening, but major differences in evolution come through different topographic settings. An ice rise is known to be located in its front (TINTO et al. 2011). Its size and orientation along the seaward boundary of the ice shelf contributes to comparably greater stability of this front. It also indicates that the buttressing effect is still in process and resists the flow of the eastern part of the grounded proportion of Thwaites Glacier (ALLEY et al. 2021; FÜRST et al. 2016). Especially in phases of accelerated ice flow, this can lead to backstress. This can be assumed to be the case in the mid 2000s. At that time, both parts of TIS were coupled. Accordingly, the advance following the huge break-up at the ice tongue also effected the eastern part. The resulting enhanced backstress may explain the steady advance of the lateral margin given that the forward advance is prevented by the topography. The discussed ongoing rifting between the two parts of TIS suggests that the coupling is gradually reduced. A closer look to the geometry of the time series and the SAR scenes gives rise to the assumption that simultaneously to the disintegration of the ice tongue, flow patterns at the eastern part changed. Close to the seaward pinning, new internal rifting indicates

concentration of stress in a frontal shear zone. This progression could be taken into consideration as a possible trigger for the lateral calving event detected in 2009.

6 Conclusion and Outlook

Considering the importance of ice shelf buttressing for sea level rise projections, the objective of this study was to create long time series of calving front dynamics. This was shown to be feasible by applying a modified version of the deep learning architecture HED-UNet to the entire archive of ERS and Envisat scenes over Pine Island Bay. The products of annual, seasonal and monthly time series do not only confirm the suitability of the data and the approach, but also provide valuable insights on ice shelf dynamics. Individual calving patterns were observed such as large break-ups (TIS) or frequent calving events (PIG). Overall, a picture of destabilisation emerges.
Future research should investigate found changes with climatic and glaciological parameters in order to deepen the process understanding. Further, an expansion of the approach to broader regions and additional sensors would be desirable. This could be a very valuable input for ice sheet models not yet considering calving mechanisms in detail.

7 References

ALLEY, K., WILD, C., LUCKMAN, A., SCAMBOS, T., TRUFFER, M., PETTIT, E., MUTO, A., WALLIN, B., KLINGER, M., SUTTERLEY, T., CHILD, S., HULEN, C., LENAERTS, J., MACLENNAN, M., KEENAN, E., & DUNMIRE, D., 2021: Two decades of dynamic change and progressive destabilization on the Thwaites Eastern Ice Shelf. The Cryosphere, **15**, 5187-5203.

BAUMHOER, C., DIETZ, A., DECH, S. & KUENZER, C., 2018: Remote Sensing of Antarctic Glacier and Ice-Shelf Front Dynamics - A Review. Remote Sensing, **10**(9), 1445.

BAUMHOER, C., DIETZ, A., KNEISEL, C. & KUENZER, C., 2019: Automated Extraction of Antarctic Glacier and Ice Shelf Fronts from Sentinel-1 Imagery Using Deep Learning. Remote Sensing, **11**(21), 2529.

BAUMHOER, C., DIETZ, A., HEIDLER, K. & KUENZER, C., 2023: IceLines – A New Data Set of Antarctic Ice Shelf Front Positions. Sci Data, **10**(138).

BINDSCHADLER, R. A., 2002: History of lower Pine Island Glacier, West Antarctica, from Landsat imagery. Journal of Glaciology, **48**(163), 536-544.

FÜRST, J., DURAND, G., GILLET-CHAULET, F., TAVARD, L., RANKL, M., BRAUN, M. & GAGLIARDINI, O., 2016: The safety band of Antarctic ice shelves. Nature Climate Change, **6**, 479-482.

HAN, H., IM, J. & KIM, H.-C., 2016: Variations in ice velocities of Pine Island Glacier Ice Shelf evaluated using multispectral image matching of Landsat time series data. Remote Sensing of Environment, **168**, 358-371.

HEIDLER, K., MOU, L., BAUMHOER, C., DIETZ, A. & ZHU, W., 2021: HED-UNet: Combined Segmentation and Edge Detection for Monitoring the Antarctic Coastline. IEEE Transactions on Geoscience and Remote Sensing, **60**, 1-14.

IMBIE., 2018: Mass balance of the Antarctic Ice Sheet from 1992 to 2017. Nature, **558**, 219-222.

LHERMITTE, S., SUN, S., SHUMAN, C., WOUTERS, B., PATTYN, F., WUITE, J., BERTHIER, E. & NAGLER, T., 2020: Damage accelerates ice shelf instability and mass loss in Amundsen Sea Embayment. Proceedings of the National Academy of Sciences, **117**(40), 24735-24741.

MACGREGOR, J.A., CATANIA, G.A., MARKOWSKI, M.S. & ANDREWS, A., 2012: Widespread rifting and retreat of ice-shelf margins in the eastern Amundsen Sea Embayment between 1972 and 2011. Journal of Glaciology, **58**(209), 458-466.

MILES, B. W. J., STOKES, C. R., JENKINS, A., JORDAN, J. R., JAMIESON, S. S. R. & GUDMUNDSSON, G. H., 2020: Intermittent structural weakening and acceleration of the Thwaites Glacier Tongue between 2000 and 2018. Journal of Glaciology, **66**(257), 485-495.

PERIYASAMY, M., DAVARI, A., SEEHAUS, T., BRAUN, M., MAIER, A. & CHRISTLEIN, V., 2022: 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.

RIGNOT, E., 2002: Ice-shelf changes in Pine Island Bay, Antarctica, 1947-2000. Journal of Glaciology, **48**(161), 247-256.

RONNEBERGER, O., FISCHER, P. & BROX, T., 2015: U-Net: Convolutional Networks for Biomedical Image Segmentation. International Conference on Medical image computing and computer-assisted intervention (MICCAI), LNCS, **9351**, 234-241.

TINTO, K. J. & BELL, R. E., 2011: Progressive unpinning of Thwaites Glacier from newly identified offshore ridge: Constraints from aerogravity. Geophysical Research Letters, **38**(20), 1-6.

WILD, C. T., ALLEY, K. E., MUTO, A., TRUFFER, M., SCAMBOS, T. A. & PETTIT, E. C., 2022: Weakening of the pinning point buttressing Thwaites Glacier, West Antarctica. The Cryosphere, **16**, 397-417.

XIE, S. & TU, Z., 2018: Holistically-Nested Edge Detection. IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 1395-1403.