The Cryosphere

SAR deep learning sea ice retrieval trained with airborne laser scanner measurements from the MOSAiC expedition

Karl Kortum[1,2], **Suman Singha**[3,4], **Gunnar Spreen**[2], **Nils Hutter**[5], **Arttu Jutila**[6,5], **and Christian Haas**[5]

[1]Remote Sensing Technology Institute, German Aerospace Center (DLR), Bremen, Germany
[2]Institute of Environmental Physics, University of Bremen, Bremen, Germany
[3]National Center for Climate Research, Danish Meteorological Institute (DMI), Copenhagen, Denmark
[4]Department of Geography, University of Calgary, Calgary, Canada
[5]Sea Ice Physics, Alfred-Wegener-Institut (AWI), Bremerhaven, Germany
[6]Finnish Meteorological Institute (FMI), Helsinki, Finland

Correspondence: Karl Kortum (karl.kortum@dlr.de)

Abstract. Automated sea ice charting from synthetic aperture radar (SAR) has been researched for more than a decade, and we are still not close to unlocking the full potential of automated solutions in terms of resolution and accuracy. The central complications arise from ground truth data not being readily available in the polar regions. In this paper, we build a data set from 20 near-coincident x-band SAR acquisitions and as many airborne laser scanner (ALS) measurements from the Multidisciplinary drifting Observatory for the Study of Arctic Climate (MOSAiC), between October and May. This data set is then used to assess the accuracy and robustness of five machine-learning-based approaches by deriving classes from the freeboard, surface roughness (standard deviation at 0.5 m correlation length) and reflectance. It is shown that there is only a weak correlation between the radar backscatter and the sea ice topography. Accuracies between 44 % and 66 % and robustness between 71 % and 83 % give a realistic insight into the performance of modern convolutional neural network architectures across a range of ice conditions over 8 months. It also marks the first time algorithms have been trained entirely with labels from coincident measurements, allowing for a probabilistic class retrieval. The results show that segmentation models able to learn from the class distribution perform significantly better than pixel-wise classification approaches by nearly 20 % accuracy on average.

1 Introduction

Sea ice classification from remote sensing and especially synthetic aperture radar (SAR) instruments has been used for monitoring the Arctic sea ice for multiple decades, with automation being proposed as early as the mid-eighties by Fily and Rothrock (1986). However, even with the inception of advanced machine learning methods and modern data analysis, a universally reliable classifier does not yet exist to retrieve sea ice classes from radar imagery. The potential for such a classifier is obvious, as humans are not able to match the speed and precision of an automated algorithm. Until now, however, this potential has yet to be fully unlocked. Human-generated ice charts (see the World Meteorological Organization overview by JCOMM, 2017) are still dominant in operational usage, despite the considerable amount of research that has been focused on the subject. These products unfortunately can provide only coarse approximate labels of the sea ice. For cross-cutting research, a more detailed and higher resolution of classification would be preferable and should be possible given the spatial resolution of the SAR sensors. As a human analyst generating an ice chart has only limited time to annotate a SAR scene, such high-resolution labels are not contained in the ice charts. Leads, for example, hot spots of ocean and atmosphere interaction (and thus of particular interest for the energy budgets), are generally not labelled in operational ice charting. At this point, with many different classifiers having been proposed and developed, in-
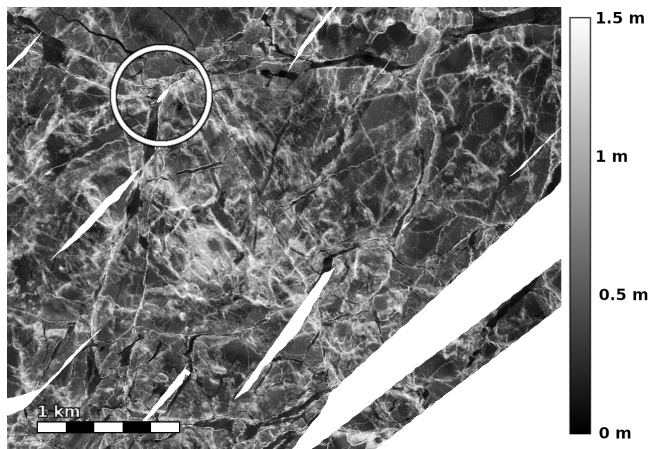
Figure 1. Section of airborne-laser-scanner-measured (ALS-measured) freeboard over the MOSAiC floe on 8 April 2020. RV *Polarstern* can be seen in the centre of the white circle. Brighter values correspond to higher freeboard values, whereas white areas indicate no data. The displayed freeboard range is 0 to 1.5 m.

cluding domain-knowledge-based and centre-pixel as well as semantic segmentation models (e.g. Kwok et al., 1992; Soh and Tsatsoulis, 1999; Hara et al., 1995; Karvonen, 2004; Ressel et al., 2015; Doulgeris, 2015; Johansson et al., 2020; Lohse et al., 2021), one must ask the question of why no meaningful direction has yet established itself in the ongoing research. The answer to the question – aside from the complexity of the subject – is two-fold. Firstly and most important is the state of the data. Although we have a great wealth of satellite SAR acquisitions of the sea ice in diverse states and conditions, we lack the corresponding ground truth information. Secondly, the constantly varying and difficult-to-predict drift and deformation of sea ice make it nearly impossible to image the same area of sea ice over longer time series to evaluate any proposed classifiers' robustness. The latter is particularly true for high-resolution imagery. These two shortcomings open the development of classifiers performing at or near the resolution of the SAR imagery to a plethora of different challenges because we have almost no way to test, iterate and improve retrieval algorithms in a structured manner. This stifles the rate at which progress in the field can be made or even recognised.

On a mission to fill gaps in our knowledge about the Arctic sea ice and its climatology, the MOSAiC expedition launched in the autumn of 2019 and the ship *Polarstern* spent a year adrift with the ice pack. Aboard, interdisciplinary teams of scientists worked to collect as much data as possible, which will help to further our understanding of one of Earth's most remote regions. With the mission came the unique opportunity to collect exactly the type of ground truth information over a long time period that is needed to test sea ice retrieval algorithms with satellite-borne SAR data being

acquired at the same time. An overview of the snow-and ice-related activities is given in Nicolaus et al. (2022).

Ice and snow transects from Itkin et al. (2021) or drilling hold the most detailed information of the underlying ice. Unfortunately, the spatial extent of these measurements is too sparse to be used for comparison with the satellite acquisitions. Aerial measurements taken from helicopters, such as the airborne laser scanner (ALS) data products by Hutter et al. (2022, 2023) used in this approach (Fig. 1), provide information about the height of the snow and/or ice surface above the local sea level (i.e. freeboard) and surface reflectance at scales of kilometres to tens of kilometres. These data are therefore prime candidates to extract ground truth information for ice classification based on roughness and thickness. Because of the efforts made during the MOSAiC expedition and the subsequent collocation for this work, the data set used here is far larger than any previously used data derived from measurements. However, it still suffers from a loss of generality from being constrained to certain regions. It is nonetheless most likely the most complete (as in large) collocated data set that we will be able to synthesise, at least until another expedition of the scope of MOSAiC comes along (which might not even happen before the first ice-free summer in the Arctic).

One prominent emerging method of segmenting image data involves machine-learning-based approaches based on convolutional neural networks, such as those published in Simonyan and Zisserman (2015), He et al. (2015), Liu et al. (2022), Ronneberger et al. (2015), Zhou et al. (2018), and Zhou et al. (2019). Advancements in the field of machine vision are being made at a rapid pace and are able to leverage improvements in chip design and the increasing amount of data that are being generated. The image-like properties of SAR acquisitions mean that this knowledge is transferable to the ice classification domain (e.g. Boulze et al., 2020; Ullah et al., 2021; Wang and Li, 2021; Kortum et al., 2021, 2022). Historically, this has been done with texture extraction and subsequent dense neural networks (Ressel et al., 2016; Singha et al., 2018; Murashkin et al., 2018), pixel-wise classification using image classifiers based on convolutional neural networks (Boulze et al., 2020; Ullah et al., 2021), and segmentation models that are able to segment an entire patch simultaneously, as detailed in Wang and Li (2021). Previous studies have also sought to use passive microwave data to derive labels (Radhakrishnan et al., 2021) which consequently concentrate on much larger scales.

In this study, we will use the unique opportunity provided by 20 instances of near-coincident (7 h time difference on average) ALS and SAR data over a period of 8 months to compare a variety of machine-learning-based classification approaches in terms of classification accuracy and robustness on classes delineated directly from measurements. For the first time, we have accurate, high-resolution sea ice topography measurements of freeboard and surface reflectance with high spatial overlap and low time differences between acqui-

sitions to truly test the capability of retrieving freeboard and (above-snow) surface-roughness-based sea ice classes from SAR data. In contrast to existing ALS and SAR data sets, such as those produced in Singha et al. (2018), the MOSAiC experiment provides the opportunity to monitor the same ice across a large temporal time span at a high resolution. The number of collocations achieved here is significantly greater than in previous studies, which enables the training of deep learning models requiring large data sets. Concretely, the questions we are trying to answer are these: how do different CNN architectures perform on labels derived from measurements (not human interpretation) and how does that influence future algorithm choices if the aim is to produce classifications near the resolution and detail of the SAR measurements?

2 Methodology

2.1 The data

The SAR component of the analysis is made up of TerraSAR-X (X-band) acquisitions in stripmap (SM) mode. The intensity scenes are normalised to σ_0 and calibration is performed as per the product specifications in Fritz et al. (2007). The resulting scenes have a pixel spacing of 3.5 m and a native radiometric resolution of 16 bits. Both horizontal send, horizontal receive (HH) and vertical send, vertical receive (VV) bands are acquired by the satellite simultaneously. This configuration of polarisations has been shown to yield valuable information for ice classification in Ressel et al. (2016) and Geldsetzer and Yackel (2009). As only two bands can be acquired simultaneously, the cross-pol band is not present in the data. Each combination of two channels will have some shortcomings, however, so this needs to be accepted. The footprint of a single scene is typically around 50×15 km. Other SAR data are not available at the resolution and frequency to enable high spatial overlap with ALS measurements (in terms of number of pixels) at small enough time differences. At higher wavelengths, especially L-band, we would expect a higher correlation between radar backscatter and ALS-derived surface roughness measured at spatial intervals of 0.5 m. This would probably translate to higher classification accuracy for deformed ice. In terms of the generality of the derived results, the complexity of the spatial distribution of classes and hence the core results derived in this study, we would expect to translate to coarser resolutions and other frequencies.

The ALS data from Hutter et al. (2023, 2022) from 20 scenes (Appendix A) between October 2019 and May 2020 are used to delineate sea ice classes. The data were acquired by flying a mowing-the-lawn pattern over the ice near the MOSAiC central observatory. The resulting ALS grid has a geospatial resolution of 0.5 m. For midwinter flights in high latitudes of $> 85°$ N, the post-processing of the helicopter IN-S/GPS data failed, and ALS data processing was performed using a lower-frequency real-time navigation solution with metre-scale undulations in GPS altitude that propagated to the surface elevation retrieved from the ALS. The undulations in the computed freeboard could be minimised using a correction calculated from swath-to-swath overlap. It should be noted that the local standard deviation of the freeboard is left intact by these processing artefacts and can still be used to derive a parameterisation of the local surface roughness where these undulations are present. An additional measurement aside from freeboard is the surface reflectance at the wavelength of the laser (1064 nm), which is useful to identify regions of young ice that have not yet been covered by snow. For the acquisitions with unphysical undulations in the freeboard measurement, freeboard was not used to delineate class labels. Instead, only classes which could be inferred from the surface roughness and reflectivity were used. The footprint of a single flight is typically around 5×10 km.

2.1.1 Collocation

For each ALS grid, the first step for collocating with SAR data is to find the SAR acquisition that is closest to the ALS measurement time, whilst still having substantial spatial overlap. Then, using the *Polarstern* ship to determine a common coordinate system, the two measurements are fused by assigning each ALS data point to the closest SAR pixel (see Kortum et al., 2021; Hendricks, 2019). In the common coordinate system, this means that the two measurements are in the same TerraSAR-X grid cell relative to the ship. Because of the difference in resolutions (0.5 m ALS and 3.5 m SAR), we obtain approximately 49 points of ALS measurements per SAR pixel. The freeboard and roughness are then computed as the respective mean and standard deviation of these points. Investigation showed that the median and mean of the local distributions were on average within less than a percent of the span of the distribution. This lends confidence that the distribution is roughly symmetrical and thus the mean and standard deviation describe the statistical nature adequately. Using the *Polarstern* as an origin of the common coordinate system is sensible, as we have accurate GPS positioning and heading to account for ice drift and rotation. The matching of the two products using this method was accurate to a couple of metres. To further improve the accuracy of collocation, final translation and rotation were then determined manually. Afterwards, the features overlapped perfectly at (TerraSAR-X) pixel resolution. The accuracy of collocation is made possible by more than daily TerraSAR-X SAR acquisitions of the MOSAiC floe, which helps keep the time differences between satellite and helicopter measurements small.

2.1.2 Determining labels

We have categorised the measured sea ice into three classes. A label is given for each SAR pixel for which ALS informa-

tion is available. For ease of reference, we give them names which are easier to contextualise. However, the exact definitions of the classes are given here. They are fully given by the ALS measurement. The three classes are open water and young ice (OW/YI), level first-year ice (LFYI), and deformed first-year and multiyear ice (DFYI/MYI). These classes we define as follows (see Fig. 2 for a visual aid):

– OW/YI is ice with a reflectance (range-corrected target echo amplitude) significantly lower than that of the surrounding snow-covered ice. Typically values around −7 dB were used as a threshold value and adjusted manually if needed. Note that finer separation here is not possible from the data alone, but we know from reports of scientists on the expedition that most ice in this class will have already formed a thin ice layer, and entirely open water was very rare during the flights.

– LFYI is snow-covered ice with a surface roughness (standard deviation of freeboard measurements at scales of the ALS grid (0.5 m) calculated over one SAR pixel ($3.5 \times 3.5\,\mathrm{m}^2$)) of less than 1 cm or a freeboard value lower than the higher inflection point in the freeboard distribution (typically around 40 cm).

– DFYI/MYI is snow-covered ice with a surface roughness of more than 10 cm or a freeboard greater than the higher inflection point in the freeboard distribution.

As detailed above, ice types are identified by thresholds in the reflectance, surface roughness or freeboard. The thresholds for the roughness and freeboard are indicated in the histograms in Fig. 2 by the different background colours. We can infer the probabilities of lying above or below a threshold for every pixel by assuming a Gaussian distribution of ALS freeboard and reflection measurements at each SAR pixel. From the 49 ALS measurements mapped to one SAR pixel, we compute the mean and standard deviation of the freeboard and can then compute the probabilities of lying below or above the globally defined freeboard thresholds using the error function. Explicitly, we integrate the area under the curve of the estimated Gaussian probability density function above and below the threshold. An example is shown in Fig. 3. Thus, we obtain "soft labels" which give the probabilities of belonging to a certain class rather than discrete classes. Assuming a Gaussian distribution allows us to also infer uncertainties in the surface roughness. One could have classified each of the 49 ALS measurements mapped to one SAR pixel and then used the relative occurrences as probabilities. This simplification to a Gaussian distribution leads to an inaccuracy of the probabilities (derived from freeboard) of only $\approx 0.16\,\%$ on average but a significantly increased computational efficiency.

The derived labels from each scene are split into three mutually exclusive connected subsets. By connected, we mean that in all but edge cases, pixels are neighbouring ones from the same subset. The training set is made up of 80 % of labels, whilst the test and validation sets consist of 10 % each. The validation data are used only to decide when to stop training. All subsets (test, training and validation) contain data from every scene. Imbalances of the classes were handled by balancing the data set for pixel-wise classifiers and weighting the classes inversely to their frequency for the segmentation approaches where an entire patch is segmented at once. Thus, the training of the networks is set up so that performing equally well for each class yields the lowest loss. As the classes are not balanced in the labels, better performance would certainly be achieved on the training data set without balancing the classes but it would hinder the generalisation of the classifier and make the results more difficult to interpret. As generalisation of a larger space of ice conditions is a property we would like to have reflected in the results as directly as possible, balancing was undertaken here.

In Fig. 4, the correlation between backscatter and surface topography measurements is shown. It becomes evident immediately that the backscatter characteristics are only very weakly correlated with the topography, and thus separation using the backscatter alone would surely be futile. This is further underlined by looking at the backscatter distributions of the delineated classes from the flight on 8 April (bottom of Fig. 2), where the correlations are relatively average in regards to all other flights. Here, it is again obvious that the backscatter characteristics are only somewhat valuable for class separation. Thus, the information needed to classify accurately must in large part be derived from contextual data.

2.2 Robustness

To test the robustness of each classifier, we will follow the same steps outlined in Kortum et al. (2022). In brief, using the *Polarstern* as an origin, a 3 km × 3 km region around the ship is used as the robustness test set. This area has been identified in 162 TerraSAR-X (TSX) SM scenes from different days. The robustness is then defined as the probability of each pixel being classified the same as in the previous and subsequent acquisitions (time between acquisitions is typically 1 d). Taking into account that the surface conditions change over time and that *Polarstern* was not perfectly stationary, this approximation of the robustness will serve only as a lower bound of the actual robustness of the classifier. In summary, we are operating under the assumption that in a time period of 2 d, the percentage of ice that has changed class (e.g. through deformation) is significantly smaller than the percentage of ice that has remained in the same class. Note that this test is only sensible for the two solid ice classes and not for the OW/YI class, which is too dynamic on a daily timescale to be analysed in this manner. The robustness is first computed for the two classes and their average is used as an indicator for the network's robustness.
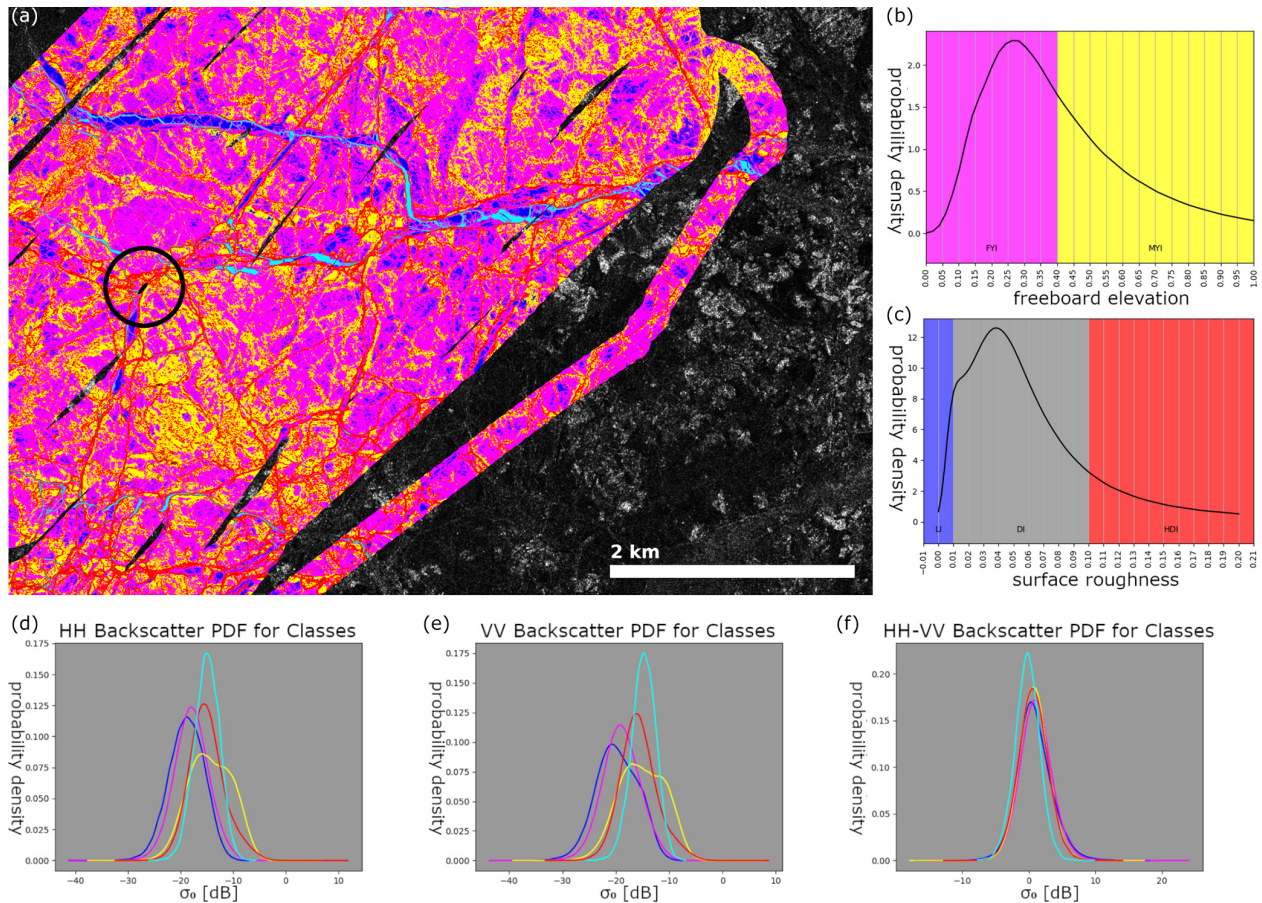
Figure 2. Derived labels from the ALS acquisition on 8 April 2020 overlaid on the HH channel of the near-coincident SAR measurement (left) and estimated probability density functions from the distributions of freeboard and surface roughness (in this case this is the local standard deviation of the freeboard) (right). Yellow areas indicate ice with a higher freeboard than the high inflection point of the distribution. Magenta represents ice with a lower freeboard than that. Red represents areas with higher surface roughness than 10 cm. Blue areas are ice with surface roughness of less than 1 cm. Cyan areas have reflectivity, indicating no snow cover (less than -7 dB echo amplitude). For this study, yellow and red as well as magenta and blue classes are combined. The grey background of the surface roughness distribution denotes the region that was not used to identify ice classes, as there was considerable mixing in this parameter region. At the bottom, approximate probability density functions (PDFs) for the σ_0 backscatter of each class across the different polarisation configurations are shown. Note that no two classes can be reliably separated using backscatter alone. All of the PDFs have been smoothed with Gaussian kernel smoothing.

2.3 The network architectures

In this paper, we will compare five different architectures: two established image classifiers in the VGG16 developed by Simonyan and Zisserman (2015) (ice classification in Khaleghian et al., 2021b) and the ConvNeXt network proposed by Liu et al. (2022) (an improvement over ResNet that is used for SAR sea ice classification in Song et al., 2021), a custom CNN (cCNN) pixel-wise classifier by Kortum et al. (2022) specifically designed for ice classification and two established segmentation models in the Unet by Ronneberger et al. (2015) (SAR sea ice classification in Nagi et al., 2021; Ren et al., 2022), and Unet++ proposed in Zhou et al. (2018, 2019) (used in Murashkin and Frost, 2021). These first three (VGG16, ConvNeXt and cCNN) and

last two (Unet and Unet++) models have one fundamental difference in that classification approaches (e.g. VGG16) are given a patch and are then asked to predict the class of the centre of the image. Segmentation approaches (e.g. Unet) are tasked to produce a label for every pixel in the patch at the same time. The exact specifications of all the models can be found in the Appendix. A short overview of the core features of the models is given here.

The oldest of the models, the VGG16, was originally developed for image classification. It uses convolutions of a filter size of three and max-pooling layers to reduce the spatial dimensions. The model architecture is completed with fully connected layers. The core idea is to extract increasingly complex features in the convolutional blocks and only keep the most prominent ones in the max pool operations. Fi-
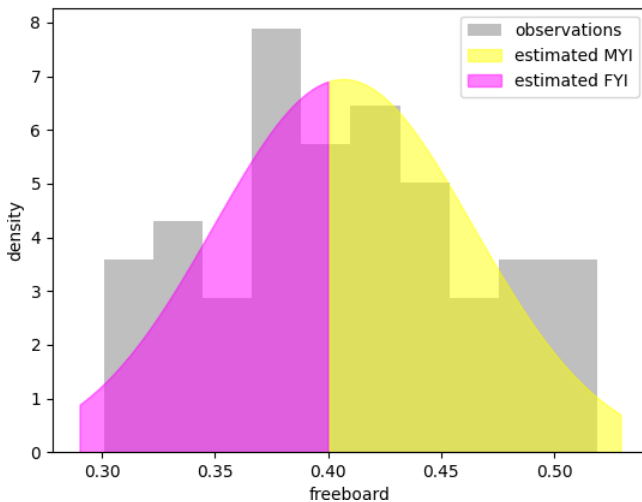
Figure 3. Soft labels are derived for one SAR pixel by assuming a Gaussian distribution (coloured) of the 49 ALS observations (grey histogram) inside of it and then integrating the area under the PDF curve above and below the threshold. In the given example, the probabilities are close to 50 %.

nally, this generates a feature vector of complex spatial features which is used by the fully connected layers to infer a class.

The ConvNeXt model is an implementation of some core advantages that self-attention-based transformer models have brought to the image classification domain in a convolutional framework. The model uses skip connections which were made a staple in large networks by the ResNet architecture and convolutional blocks using large filters, an inverted bottleneck and depth-wise convolutions inspired by transformer models. It is a more modern design that can achieve significantly higher scores on image classification tasks yet is not designed for centre-pixel classification as is typical in sea ice retrieval applications.

The custom CNN model uses multiple zoom levels as inputs and is constrained with fewer parameters and fewer layers than the other model. The central design philosophy here is to avoid overfitting through model constraint. This model architecture was optimised with performance on coarser human annotations in mind.

The Unet model follows a similar approach to the VGG16 in the initial (encoding) stage of the network where features are disseminated by convolutional layers. However, in the second part of the model, these features are upsampled again to create a 2D semantic segmentation map. The core advantage this has over centre-pixel classifiers lies in the fact that inter-label dependencies and relations can be learned and exploited by the architecture.

The Unet++ keeps the encoding–decoding framework of the original Unet but adds more intermediate layers at various scales and fuses the features from multiple scales to make a more informed prediction. We chose to average over the

features of the multiple output layers in the deep supervision part of the model.

It is important to note that as little as possible was changed about the architectures themselves to keep true to the models that were proposed in the original publications – otherwise the results would be harder to interpret with the efforts and successes of optimisation techniques being an unknown factor.

2.4 Training

During training, the networks are tasked with minimising the categorical cross-entropy between the output and the label distributions. This allows us to fit the probabilities of each class occurring at each pixel which we can infer from the ALS measurements. Minimising the cross-entropy gives the same result as minimising the Kullback–Leibler divergence (KLD). As this serves as a benchmark and comparison of these models concerning their applicability for sea ice retrieval, no further optimisations have taken place. For each of the model architectures, 10 separate instances are from scratch. Training is stopped using a spatially independent validation set (10 % of data) after the model does not improve in the last 100 000 training samples. Testing is done on another spatially non-overlapping 10 % of data and the remaining (disjoint) 80 % of data are used for training (see Fig. 5 for details). Training multiple instances allows some additional insight into the variabilities. The ingested SAR data are preprocessed by converting each band σ_0 and then applying a logarithm. The incidence angle is provided in a third channel. The Adam optimizer proposed by Kingma and Ba (2017) is used to train the models with the learning rate set to 10^{-4}. The size of each patch to be classified is 256×256 pixels, except for the cCNN model, which receives input patches at various scales (a 5×5, a 16×16 and a 64×64 pixel patch).

3 Results

The performance of different network architectures can be seen in Table 1. They paint a clear picture of segmentation models' (Unet and Unet++) improvement over centre-pixel classification approaches. Of the pixel-wise classification approaches, the custom CNN classifier performed best on unseen test data, yet it was still significantly inferior to the segmentation models. We speculate that part of the reason for this is the high spatial resolution of the labels, as we get a label for every pixel from most of the ALS measurements. The pixel-wise classifiers cannot make use of any relationships between spatial properties of labels like shape, sparsity and correlations. This seems to be detrimental to their performance. Except for the cCNN model, which was designed to avoid overfitting, the other centre-pixel classifiers show a large discrepancy between training and test scores, whilst the
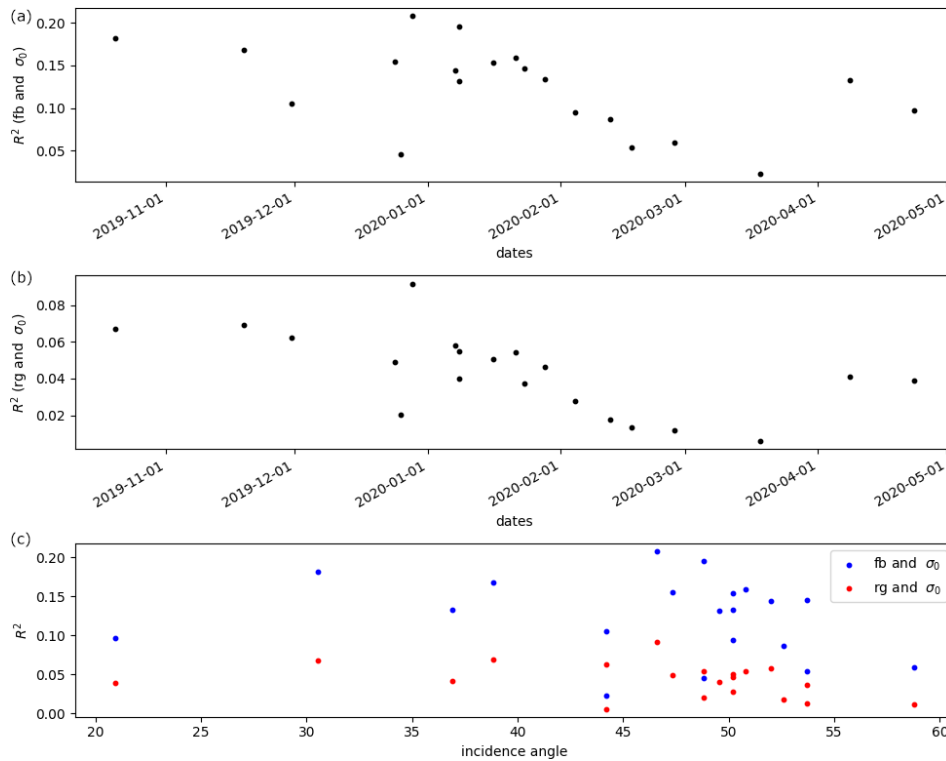
Figure 4. Evolution of correlations between freeboard [fb] or surface roughness [rg] and HH SAR backscatter [σ_0] over time. Panel **(c)** shows correlations plotted against incidence angle. Note that the surface roughness is measured at the snow–atmosphere interface and at correlation lengths of 0.5 m, whilst the SAR sensor is most sensitive to the ice–snow interface and roughness at correlation lengths at the wavelength of the sensor, which is only 3.1 cm. The same analysis with VV channels gives very similar results.
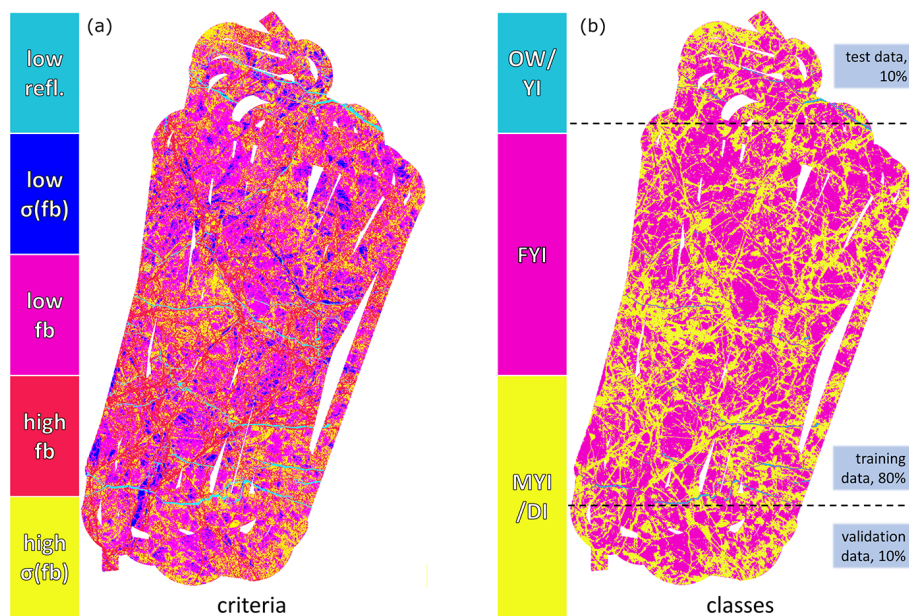

Figure 5. Split of training, test and validation data demonstrated on a helicopter product from 23 April. The left-hand side shows the different criteria (as in Fig. 2) and the right-hand side shows the derived training data with splits indicated by the dotted lines. The same procedure was carried out for all 20 flights.

semantic segmentations models have generalised much more effectively.

A more detailed analysis of the output of different models (Fig. 6) shows how the VGG16 and ConvNeXt models struggle to relate all the information of the patch to only the classification of the central pixel, leading to a diffuse-looking classified scene. This seems most pronounced for the ConvNeXt model. A possible reason for this is that the larger convolutional kernels (7×7 in contrast to 3×3) used in the architecture. (A retraining with a kernel size of 3×3 confirmed this increased the average accuracy by around 5 %.) The cCNN model seems to struggle with using contextual data to separate rough ice and young ice. In general, the predicted probabilities at each pixel are higher in the non-dominant class, leading to a seemingly different colour palette in this visualisation. The Unet and Unet++ classifications are largely similar. Some difficulty in the separation of deformed and young-ice signatures persists as can be seen in the mixing of yellow and cyan areas.

It is also worth pointing out that the very same cCNN model and a VGG16 model performed at accuracies around 85 %–95 % on manual labels in Kortum et al. (2022), illustrating the difference between training and testing on quantitatively measured labels in contrast to human-generated annotations. In Ren et al. (2022), the Unet model is reported to perform sea ice and open-water separation on manual labels at 93 %–95 % accuracy. Wang and Li (2021) report accuracies of 96 % for the same task, using ice charts as training data and test data, and 94 % accuracy when comparing to an operational sea ice cover product (Interactive Multisensor Snow and Ice Mapping System, IMS). Murashkin et al. (2018) show classification accuracies of the Unet++ model around 96 % on manually labelled training and test data across six classes.

Whilst the mean KLDs are in accordance with the accuracies, the spread (SD) of the KLDs across the model populations seems to be very similar across all models, and there is no clear gap between segmentation and classification approaches. Overall, we cannot say that one model converges more reliably than another as would be suggested by the accuracies alone. It is also apparent that the cCNN model does not perform well in the robustness scores on this data set. This model is considerably smaller than the others (in terms of parameter count) and was heavily optimised using a different data set, which seems to have come at the cost of flexibility/generality of the architecture. The spread of the robustness of the segmentation models seems to be considerably smaller that those of the centre-pixel classifier models – additionally indicating these approaches are more reliable for ice classification from SAR.

The classifications (e.g. in Fig. 7) show a very plausible set of results that align with the observations of members on board the expedition. The fine labels at a high resolution seem to have transcended into a similarly detailed classification map. The examples in Fig. 8 also illustrate a general increase in deformation in the first-year ice. For instance, the magenta FYI area close to *Polarstern*, marked by a black square in Fig. 7, is becoming progressively more deformed (detail in Fig. 8). The areas most prone to error seem to be the OW/YI classifications. This is to be expected, as they are naturally the most sparse in the training data set. Additionally, they are very dynamic, which leads to extremely diverse backscatter properties that can be exhibited, in turn making them more difficult to classify.

We also observe a decreasing correlation between backscatter and surface topography variables from the onset of the expedition until early April – particularly during January and February (Fig. 4), where the MOSAiC expedition was met by numerous storms. Some of the decorrelation can be accounted for because of snow accumulation and redistribution, but it is difficult to quantify this phenomenon. However, the incidence angles of the scenes also change, and the increase in incidence angle beyond 45° is shown to lead to a continuous decrease in correlation between ice surface characteristics and SAR backscatter.

It should also be noted that the use of one-hot-encoded labels leads to a decrease in accuracy of 8 % in comparison to smooth/probabilistic labels for the Unet architecture.

4 Discussion

The top models in our investigation perform at around 68 % accuracy on the test data set (Table 1). The predictions of the segmentation models are approximately 20 % more accurate than those of the classification models. The only concrete difference between these models is that the segmentation approaches can learn from the distribution of labels, which appears to be highly important to generalise unknown regions. The centre-pixel classifiers show a much larger difference between test and training sets. Even the highest accuracies measured here are considerably lower than what many authors report for algorithms trained with human-made labels. To understand these discrepancies, we will discuss the main differences between these measured labels and human annotation.

The measured labels used in this study have some underlying difficulties because we do not know the snow depth and density or how strong the correlation of freeboard and ice thickness is, and we cannot eliminate this error. Also, the reflectance used to disseminate young ice and open water is based purely on the coverage of the surface being snow-free and thus not directly correlated with ice age. For example, if thin ice has formed, the atmospheric conditions will dictate whether or not snow has gathered on top or if the bare ice is visible to the sensor. Thus, the quality of labels could still be improved upon if more information were available.

To assess the impact of the individual thresholds (e.g. the location of the inflection point in the freeboard distribution), we also evaluated the top-performing Unet architecture in

Table 1. Network performances on the independent test set after training. For brevity, we shortened accuracy to "acc" and robustness to "rb". The means and standard deviations are computed from the 10 models in the population for every architecture. Best-in-category results on independent test sets are highlighted in bold font. For every model, 10 instances were trained. The Unet and Unet++ architectures show significantly better performance than the other models tested.

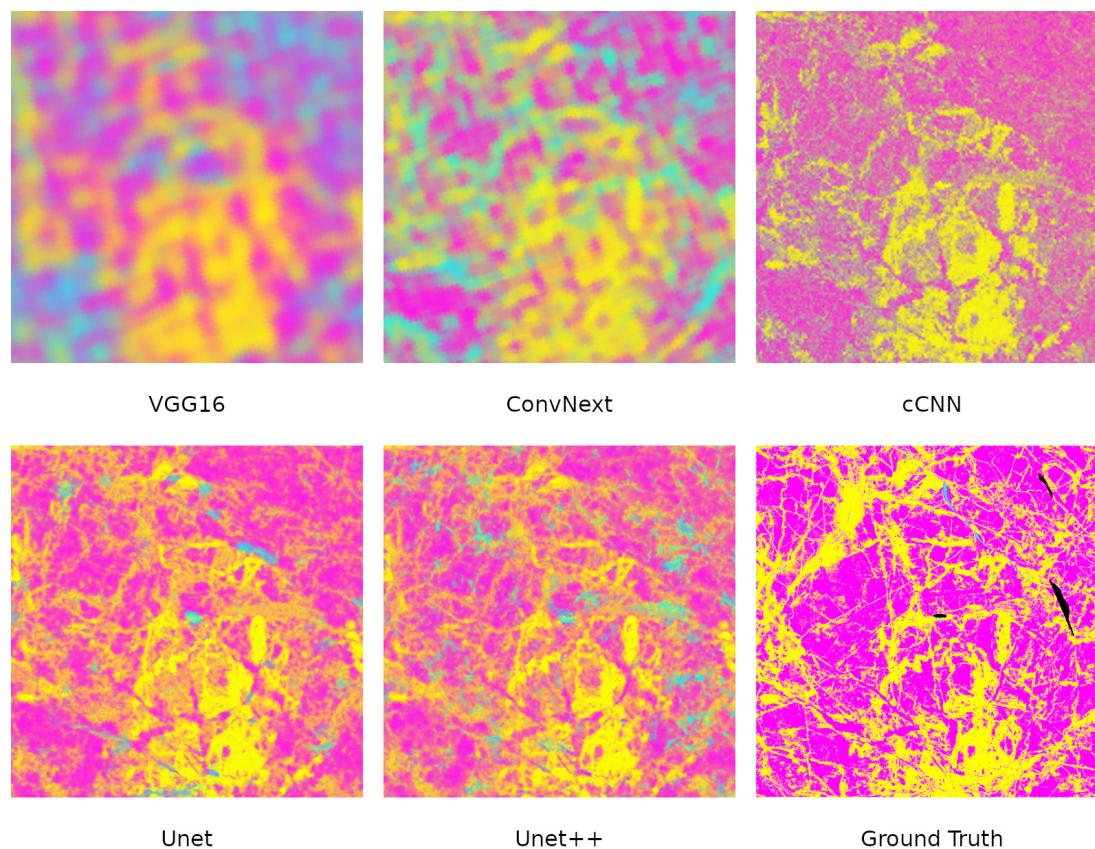| Model	train acc. [%]	test acc [%]	train KLD	test KLD	rb. [%]
VGG16	57.80 ± 7.07	44.00 ± 1.79	0.6468 ± 0.0873	0.838 ± 0.0416	80.31 ± 5.59
ConvNeXt	56.63 ± 5.21	44.72 ± 1.47	0.704 ± 0.0742	0.8784 ± 0.0368	81.45 ± 5.6
cCNN	47.14 ± 1.92	47.65 ± 1.85	0.7635 ± 0.0188	0.7724 ± 0.0269	71.12 ± 15.85
Unet	67.84 ± 2.39	65.18 ± 1.08	0.5769 ± 0.0307	$\mathbf{0.6057 \pm 0.0316}$	$\mathbf{83.12 \pm 1.96}$
Unet++	66.06 ± 1.6	$\mathbf{66.22 \pm 1.3}$	0.6104 ± 0.0201	0.6374 ± 0.022	82.82 ± 1.33

Figure 6. Comparison of classifications from different models randomly selected from the 10 instances trained. Colours are the same as the classes discussed above, but the intensity is given by the predicted probabilities, so mixed colours can occur. The scene was acquired over the *Polarstern* (centre of the images) on 14 January. The false-colour composition consists of HH, VV and HH/VV channels, normalised with a tanh function. The area shown is 6 by $6 \, \text{km}^2$.

the same data set but excluded points near the thresholds. To do this, we did not consider labels where the certainty in the most probable class was lower than 90 %. For example, regions with a local standard deviation of approximately 3 cm, which means points within 6 cm of the thresholds are not considered and the exact value of the thresholds have little bearing on the data considered. In the case of the test data set, these data points account for 24.1 % of all data. Under these circumstances, the average accuracy of the Unet model is 72.5 %, which is an increase of only 4.18 %, although 24.1 % of the least-certain labels were removed. Thus, we can conclude that the exact location of the thresholds had only marginal impact on model performance, lending increased confidence that the model performances are representative of performances evaluated against ground truth.
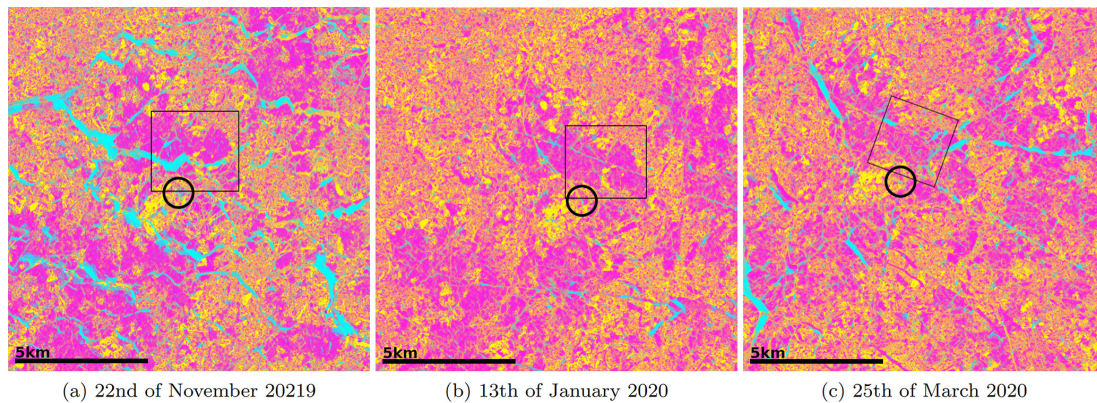
(a) 22nd of November 20219   (b) 13th of January 2020   (c) 25th of March 2020

Figure 7. Collection of classified subscenes (Unet, pixel spacing = 3.5 m) including the MOSAiC floe, after a storm **(a)**, in calm conditions with some shearing indications **(b)** and with some breakup of the ice cover visible **(c)**. The *Polarstern* location is indicated by the black circle. The DFYI/MYI class probability is displayed in yellow, the LFYI probability in magenta and the OW/YI probability in cyan. The black square marks the area shown at full resolution below (Fig. 8)

(a) 22nd of November 2019 - DFYI: 3.1%   (b) 13th of January 2020 - DFYI: 5.6%   (c) 25th of March 2020 - DFYI: 13.5%
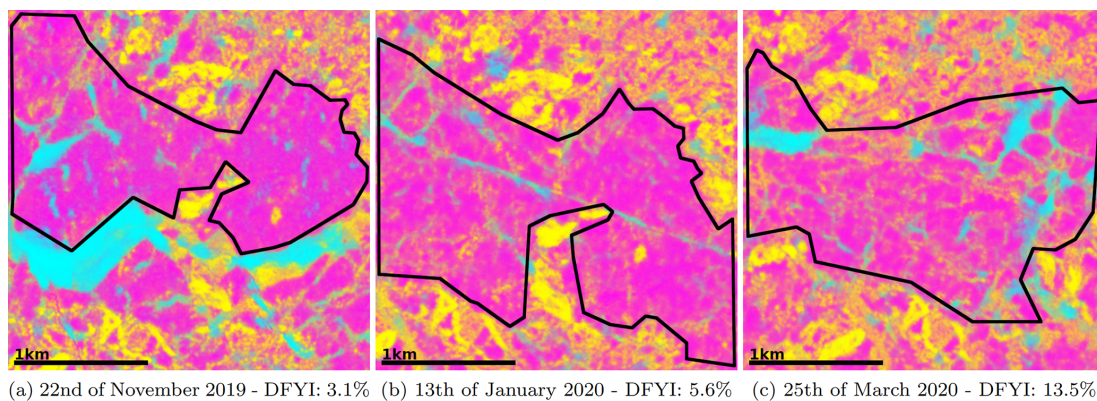
Figure 8. Full-resolution excerpt from the scenes show in Fig. 7. The classified images reflect the increased deformation of the first-year ice area over time accurately as the DFYI occurrence rises. The DFYI fraction is computed inside of the black border. In the first scene, some misclassification of the open lead (cyan) as older, deformed ice (yellow) is seen (outside of the area we are computing the DFYI fraction in). This is a common issue in SAR sea ice retrieval, as the backscatter can become very similar (Guo et al., 2023).

When comparing the data presented here with human annotations/ice charts, one must mention the resolution. In our case, every individual pixel gets its own class and there is no semantic grouping of pixels into the same class based on proximity or likeness. This is in stark contrast to ice charts, where the labels are made up of only a few polygons per scene. Even when not training from such ice charts, humans generating training data for algorithms at high resolutions generally limit themselves to areas which they can confidently identify. Not much can be said about the correctness of these labels per se, but one should keep in mind that in these instances, the accuracy achieved by the classifier is constrained to those easy-to-identify regions and are therefore not representative of the classifier's performance on the whole. Because of the size of SAR acquisitions, obtaining labels at pixel resolution from human annotation is not feasible. The great advantage that labels from measurements have

is that they are truly indicative of performance on the entire scope of ice conditions in the scene (every pixel is labelled, and thus there is no selection bias). Only by holding the testing of our high-resolution retrieval algorithms to this standard can we show with certainty when an improved method of classification is developed, but we are lacking available data sets to do so.

This study had only a small effective study region and a large temporal span to test the diverse conditions. Overall, the constancy of the ice in the scenes should only improve the classifiers' performances. Unfortunately, the 20 helicopter flights are not quite enough to make meaningful statements about temporal changes in performance, as the differences in performance will be outweighed by the local conditions in the scene. Additionally, seasons in the data where one would expect the classification to be most difficult (freeze-up and pre-early-melt onset) are only very sparsely represented in

the data. This means the contributions of the data sparsity, seasonality and spatial variability cannot be meaningfully separated.

In the summer season, the ice surface is dominated by wet snow, bare ice and melt ponds, and more open water is found between floes. The spatial distribution of classes is very distinctive between the surface types, so one can expect the main result of the difference between centre-pixel classifiers and segmentation models to persist.

In most data-driven approaches to classification, the performance of the classifier is limited by the quality of the labels. Therefore, one should be careful when using manually labelled data, such as ice charts, as ground truth. These practices are common in the current research, as not many other sources of labels are available. However, the potential is much greater than that. Of course, the great challenge remains that high-resolution measurements are very sparse.

Because the MOSAiC mission provided us with an unmatched opportunity for training and testing algorithms with measured labels over a long time period, this study has made obvious that there is considerable room for improvement even with modern deep learning algorithms. It needs to be mentioned that due to the spatial constraint to the area near the MOSAiC floe, the training data set does not capture the full extent of possible winter ice conditions in the Arctic; thus, we cannot expect the classifier to perform equally well on a pan-Arctic scale. Instances of OW/YI are very sparse, and their entire span of possible conditions and consequent radar responses are not covered well by data. Since a better in situ data set is probably not going to emerge in the near future, it is clear that measured labels alone are not enough to train a stable algorithm that can deal with the full span of ice conditions. It seems that to achieve this, one would need to leverage a great number of scenes without labels. Semi-supervised and self-supervised approaches come to mind. Some first examples of their development exist for optical data (Han et al., 2019), ice and open-water discrimination from SAR (Li et al., 2015; Khaleghian et al., 2021a), and sea ice classes from SAR (Imber, 2022).

5 Conclusions

The MOSAiC expedition enabled the generation of a large data set (ca. 20 million data points) of SAR acquisitions and appropriate labels delineated from in situ laser scanning measurements. It has become clear that both the freeboard and the above-snow surface roughness (at correlation lengths of 50 cm) are only weakly correlated with x-band SAR backscatter, with average R^2 values of 0.124 and 0.043, respectively. We have shown that deep-learning-based segmentation approaches such as the Unet can approximate these labels from the SAR measurement at accuracies around 68 %. We measured the performance of modern network architectures on a representative set of labels for the first time,

and it is much more difficult to classify at this high detail than at coarser-label resolution (e.g. ice charts). From the performances of the different models, we can conclude that the semantic segmentation approaches' advantage of being able to make use of the spatial relationship of predictions is crucial (20 % accuracy) to the generalising of unseen regions. It is notable that these label distributions at the scale of the measurement resolution are not contained in ice charts or human annotations, which suggests that accurate classification at the resolution of the SAR measurement when it is trained on human-annotated labels is improbable. As a more comprehensive data set than created here is unlikely to be acquired in the near future, newly developed classifiers aiming to classify at the resolution of the sensor will need to find some way to gain access to the spatial ice-type distributions to be successful.

Appendix A: List of Helicopter Flights

Table A1. List of the 20 helicopter flights used in this research. Data are published in Hutter et al. (2022).

1	20191020_01_PS122-1_2-167
2	20191119_01_PS122-1_8-23
3	20191130_01_PS122-1_9-98
4	20191224_01_PS122-2_17-98
5	20191225_01_PS122-2_17-99
6	20191228_01_PS122-2_17-101
7	20200107_01_PS122-2_19-44
8	20200108_01_PS122-2_19-46
9	20200108_03_PS122-2_19-52
10	20200116_01_PS122-2_20-52
11	20200121_01_PS122-2_21-41
12	20200123_02_PS122-2_21-78
13	20200128_01_PS122-2_22-16
14	20200204_01_PS122-2_23-14
15	20200212_01_PS122-2_24-31
16	20200217_02_PS122-2_25-8
17	20200227_01_PS122-3_29-49
18	20200318_01_PS122-3_32-42
19	20200408_01_PS122-3_35-49
20	20200423_01_PS122-3_37-63

Appendix B: Network architectures

We briefly present the network architectures used in this investigation. We make use of the following conventions to keep the figures concise. FCX is short for a fully connected layer with X neurons. ConvX x Y denotes a 2D convolutional layer with filter sizes X and number of filters Y. Unless otherwise specified, the convolutional layers have stride 1. If a layer has multiple inputs, they are concatenated before being parsed to the layer.
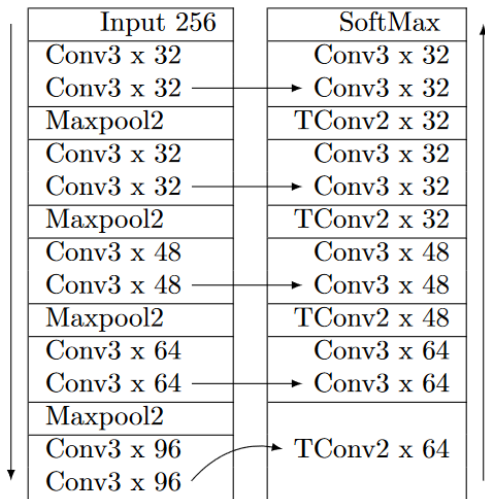
Table B1. VGG16 architecture as used in the paper. Published in Simonyan and Zisserman (2015). The ReLU activation is used throughout the network. The padding is set to "same".

Input 256
Conv3 x 64
Conv3 x 64
Maxpool2
Conv3 x 96
Conv3 x 96
Maxpool2
Conv3 x 128
Conv3 x 128
Conv3 x 128
Maxpool2
Conv3 x 192
Conv3 x 192
Conv3 x 192
Maxpool2
Conv3 x 256
Conv3 x 256
Conv1 x 256
Maxpool2
FC256
FC128
FC128
SoftMax

Table B2. The Unet architecture as used in this paper and published in Ronneberger et al. (2015). The ReLU activation is used throughout the network and the padding is set to "same" where applicable.

Input 256	SoftMax
Conv3 x 32	Conv3 x 32
Conv3 x 32 →	→ Conv3 x 32
Maxpool2	TConv2 x 32
Conv3 x 32	Conv3 x 32
Conv3 x 32 →	→ Conv3 x 32
Maxpool2	TConv2 x 32
Conv3 x 48	Conv3 x 48
Conv3 x 48 →	→ Conv3 x 48
Maxpool2	TConv2 x 48
Conv3 x 64	Conv3 x 64
Conv3 x 64 →	→ Conv3 x 64
Maxpool2	
Conv3 x 96	→ TConv2 x 64
Conv3 x 96	

Table B3. The ConvNeXt-T architecture used in this paper. Developed in Liu et al. (2022). The term "st" denotes stride.

Input 256
Conv4 (stride 4) x 96
ConvNx x 96
ConvNx x 96
ConvNx x 96
Conv2 (st 2) x 96
ConvNx x 192
ConvNx x 192
ConvNx x 192
Conv2 (st 2) x 192
ConvNx x 384
ConvNx x 384
ConvNx x 384
ConvNx x 384
ConvNx x 384
ConvNx x 384
ConvNx x 384
ConvNx x 384
ConvNx x 384
Conv2 (st 2) x 384
ConvNx x 768
ConvNx x 768
ConvNx x 768
GlobalAvgPool2D
LayerNorm
SoftMax

where:

$ConvNx \ x \ n$ =

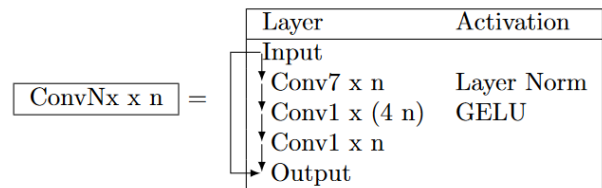| Layer	Activation
Input	
Conv7 x n	Layer Norm
Conv1 x (4 n)	GELU
Conv1 x n	
Output	

Table B4. The custom CNN architecture from Kortum et al. (2022) used in this paper. The inputs at different scales are flattened and concatenated before being output to the fully connected layers. Leaky ReLU is used for activation and padding is set to "valid". The 16×16 pixel input is downscaled from the original scene by a factor of 5 and the 64×64 pixel input is a square cutout that is rescaled so that the width of the entire scene is 64 pixels. The 1D input contains the relative coordinates of the pixel in the 64×64 pixel input. The term "st" denotes stride.

Input 64	Input 16	Input 5
Conv3 x 16	Conv3 x 16	Conv2 x 16
Conv3 (st 2) x 32	Conv3 x 16	Conv2 x 16
Conv3 (st 2) x 32	Conv3 (st 2) x 32	Conv2 x 32
Conv3 (st 2) x 32	Conv3 (st 2) x 32	
Conv3 (st 2) x 64		

Input1D 2
FC64
FC32
FC16
FC8
SoftMax
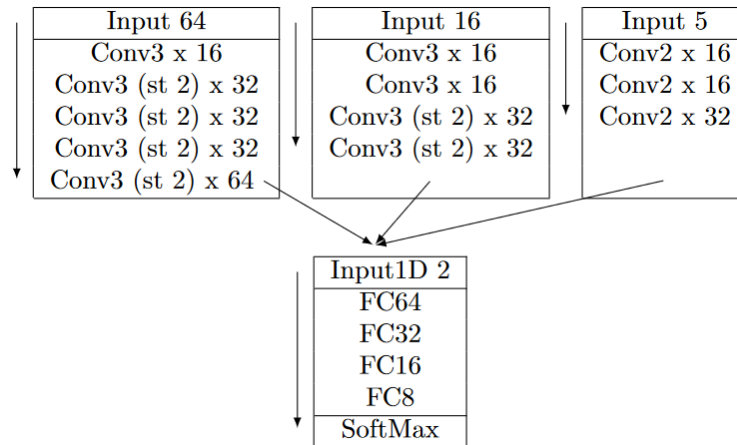
Table B5. Unet++ architecture used in this paper, published in Zhou et al. (2018, 2019). Note that the left column is identical to the downwards convolution side of the regular Unet, and the lowest rows from left to right form the upwards side of the Unet. The Unet++ then uses extra layers in between to extend the architecture. All layers within a cell are considered to be a block, so they are all executed before parsing the output to the next block. All layers marked "softmax" are averaged before the final linear layer and the softmax are applied. ReLU is used as the activation function throughout and the padding is set to "same".

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

References

Boulze, H., Korosov, A., and Brajard, J.: Classification of Sea Ice Types in Sentinel-1 SAR Data Using Convolutional Neural Networks, Remote Sens., 12, 2165, https://doi.org/10.3390/rs12132165, 2020.

Doulgeris, A. P.: An Automatic $\mathcal{U}$-Distribution and Markov Random Field Segmentation Algorithm for PalSAR Images, IEEE T. Geosci. Remote, 53, 1819–1827, https://doi.org/10.1109/TGRS.2014.2349575, 2015.

Fily, M. and Rothrock, D. A.: Extracting Sea Ice Data from Satellite SAR Imagery, IEEE T. Geosci. Remote, GE-24, 849–854, https://doi.org/10.1109/TGRS.1986.289699, 1986.

Fritz, T., Mittermayer, J., Schaettler, B., Buckreuss, S., Werninghaus, R., and Balzer, W.: Level 1b Product Format Specification, DLR: TerraSAR-X Ground Segment, https://www.intelligence-airbusds.com/files/pmedia/public/r460_9_030201_level-1b-product-format-specification_1.3.pdf (last access: November 2022), 2007.

Geldsetzer, T. and Yackel, J. J.: Sea ice type and open water discrimination using dual co-polarized C-band SAR, Can. J. Remote Sens., 35, 73–84, https://doi.org/10.5589/m08-075, 2009.

Guo, W., Itkin, P., Singha, S., Doulgeris, A. P., Johansson, M., and Spreen, G.: Sea ice classification of TerraSAR-X ScanSAR images for the MOSAiC expedition incorporating per-class incidence angle dependency of image texture, The Cryosphere, 17, 1279–1297, https://doi.org/10.5194/tc-17-1279-2023, 2023.

Han, Y., Zhao, Y., Zhang, Y., Wang, J., Yang, S., Hong, Z., and Cao, S.: A Cooperative Framework Based on Active and Semi-supervised Learning for Sea Ice Classification using EO-1 Hyperion Data, T. Jpn. Soc. Aeronaut. S., 62, 318–330, https://doi.org/10.2322/tjsass.62.318, 2019.

Hara, Y., Atkins, R., Shin, R., Kong, J. A., Yueh, S., and Kwok, R.: Application of neural networks for sea ice classification in polarimetric SAR images, IEEE T. Geosci. Remote, 33, 740–748, https://doi.org/10.1109/36.387589, 1995.

He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, arXiv [preprint], https://doi.org/10.48550/ARXIV.1512.03385, 2015.

Hendricks, S.: Ice Drift – Transformation of GPS positions into a translating and rotating coordinate reference system, https://gitlab.awi.de/floenavi-crs/icedrift (last access: October 2021), 2019.

Hutter, N., Hendricks, S., Jutila, A., Birnbaum, G., von Albedyll, L., Ricker, R., and Haas, C.: Merged Grids of Sea-Ice or snow freeboard from helicopter-borne laser scanner during the MOSAiC Expedition, version 1, PANGEA [data set], https://doi.org/10.1594/PANGAEA.950896, 2022.

Hutter, N., Hendricks, S., Jutila, A., Birnbaum, G., von Albedyll, L., Ricker, R., and Haas, C.: Digital elevation models of the sea-ice surface from airborne laser scanning during MOSAiC, Sci. Data, 10, 729, https://doi.org/10.1038/s41597-023-02565-6, 2023.

Imber, J.: Generative Network For Semi-supervised Sea Ice Classification, TechRxiv [preprint], https://doi.org/10.36227/techrxiv.21081136.v1, 2022.

Itkin, P., Webster, M., Hendricks, S., Oggier, M., Jaggi, M., Ricker, R., Arndt, S., Divine, D. V., von Albedyll, L., Raphael, I., Rohde, J., and Liston, G. E.: Magnaprobe snow and melt pond depth measurements from the 2019–2020 MOSAiC expedition, PAN-

GAEA [data set], https://doi.org/10.1594/PANGAEA.937781, 2021.

JCOMM: Sea-Ice Information Services in the World, World Meteorological Organization, https://doi.org/10.25607/OBP-1325, 2017.

Johansson, A. M., Malnes, E., Gerland, S., Cristea, A., Doulgeris, A., Divine, D., Pavlova, O., and Lauknes, T. R.: Consistent ice and open water classification combining historical synthetic aperture radar satellite images from ERS-1/2, Envisat ASAR, RADARSAT-2 and Sentinel-1A/B, Ann. Glaciol., 61, 1–11, https://doi.org/10.1017/aog.2019.52, 2020.

Karvonen, J.: Baltic Sea ice SAR segmentation and classification using modified pulse-coupled neural networks, IEEE T. Geosci. Remote, 42, 1566–1574, https://doi.org/10.1109/TGRS.2004.828179, 2004.

Khaleghian, S., Ullah, H., Kræmer, T., Eltoft, T., and Marinoni, A.: Deep Semisupervised Teacher–Student Model Based on Label Propagation for Sea Ice Classification, IEEE J. Sel. Top. Appl., 14, 10761–10772, https://doi.org/10.1109/JSTARS.2021.3119485, 2021a.

Khaleghian, S., Ullah, H., Kræmer, T., Hughes, N., Eltoft, T., and Marinoni, A.: Sea Ice Classification of SAR Imagery Based on Convolution Neural Networks, Remote Sens., 13, 1734, https://doi.org/10.3390/rs13091734, 2021b.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, arXiv [preprint], https://doi.org/10.48550/arXiv.1412.6980, 2017.

Kortum, K., Singha, S., Spreen, G., and Hendricks, S.: Automating Sea Ice Characterisation from X-Band SAR with Co-Located Airborne Laser Scanner Data Obtained During The MOSAiC Expedition, in: Geoscience and Remote Sensing Symposium, 2021 IEEE International, https://doi.org/10.1109/IGARSS47720.2021.9553340, 2021.

Kortum, K., Singha, S., and Spreen, G.: Robust Multiseasonal Ice Classification From High-Resolution X-Band SAR, IEEE T. Geosci. Remote, 60, 1–12, https://doi.org/10.1109/TGRS.2022.3144731, 2022.

Kwok, R., Rignot, E., Holt, B., and Onstott, R.: Identification of sea ice types in spaceborne synthetic aperture radar data, J. Geophys. Res.-Oceans, 97, 2391–2402, https://doi.org/10.1029/91JC02652, 1992.

Li, F., Clausi, D. A., Wang, L., and Xu, L.: A semi-supervised approach for ice-water classification using dual-polarization SAR satellite imagery, in: 2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), 28–35, https://doi.org/10.1109/CVPRW.2015.7301380, 2015.

Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., and Xie, S.: A ConvNet for the 2020s, arXiv [preprint], https://doi.org/10.48550/ARXIV.2201.03545, 2022.

Lohse, J., Doulgeris, A., and Dierking, W.: Incident Angle Dependence of Sentinel-1 Texture Features for Sea Ice Classification, Remote Sens., 13, 552, https://doi.org/10.3390/rs13040552, 2021.

Murashkin, D. and Frost, A.: Arctic Sea ICE Mapping Using Sentinel-1 SAR Scenes with a Convolutional Neural Network, in: 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, 5660–5663, https://doi.org/10.1109/IGARSS47720.2021.9553206, 2021.

Murashkin, D., Spreen, G., Huntemann, M., and Dierking, W.: Method for detection of leads from Sentinel-1 SAR images, Ann. Glaciol., 59, 1–13, https://doi.org/10.1017/aog.2018.6, 2018.

Nagi, A. S., Kumar, D., Sola, D., and Scott, K. A.: RUF: Effective Sea Ice Floe Segmentation Using End-to-End RES-UNET-CRF with Dual Loss, Remote Sens., 13, 2460, https://doi.org/10.3390/rs13132460, 2021.

Nicolaus, M., Perovich, D. K., Spreen, G., Granskog, M. A., von Albedyll, L., Angelopoulos, M., Anhaus, P., Arndt, S., Belter, H. J., Bessonov, V., Birnbaum, G., Brauchle, J., Calmer, R., Cardellach, E., Cheng, B., Clemens-Sewall, D., Dadic, R., Damm, E., de Boer, G., Demir, O., Dethloff, K., Divine, D. V., Fong, A. A., Fons, S., Frey, M. M., Fuchs, N., Gabarró, C., Gerland, S., Goessling, H. F., Gradinger, R., Haapala, J., Haas, C., Hamilton, J., Hannula, H.-R., Hendricks, S., Herber, A., Heuzé, C., Hoppmann, M., Høyland, K. V., Huntemann, M., Hutchings, J. K., Hwang, B., Itkin, P., Jacobi, H.-W., Jaggi, M., Jutila, A., Kaleschke, L., Katlein, C., Kolabutin, N., Krampe, D., Kristensen, S. S., Krumpen, T., Kurtz, N., Lampert, A., Lange, B. A., Lei, R., Light, B., Linhardt, F., Liston, G. E., Loose, B., Macfarlane, A. R., Mahmud, M., Matero, I. O., Maus, S., Morgenstern, A., Naderpour, R., Nandan, V., Niubom, A., Oggier, M., Oppelt, N., Pätzold, F., Perron, C., Petrovsky, T., Pirazzini, R., Polashenski, C., Rabe, B., Raphael, I. A., Regnery, J., Rex, M., Ricker, R., Riemann-Campe, K., Rinke, A., Rohde, J., Salganik, E., Scharien, R. K., Schiller, M., Schneebeli, M., Semmling, M., Shimanchuk, E., Shupe, M. D., Smith, M. M., Smolyanitsky, V., Sokolov, V., Stanton, T., Stroeve, J., Thielke, L., Timofeeva, A., Tonboe, R. T., Tavri, A., Tsamados, M., Wagner, D. N., Watkins, D., Webster, M., and Wendisch, M.: Overview of the MOSAiC expedition: Snow and sea ice, Elementa, 10, 000046, https://doi.org/10.1525/elementa.2021.000046, 2022.

Nixdorf, U., Dethloff, K., Rex, M., Shupe, M., Sommerfeld, A., Perovich, D. K., Nicolaus, M., Heuzé, C., Rabe, B., Loose, B., Damm, E., Gradinger, R., Fong, A., Maslowski, W., Rinke, A., Kwok, R., Spreen, G., Wendisch, M., Herber, A., Hirsekorn, M., Mohaupt, V., Frickenhaus, S., Immerz, A., Weiss-Tuider, K., König, B., Mengedoht, D., Regnery, J., Gerchow, P., Ransby, D., Krumpen, T., Morgenstern, A., Haas, C., Kanzow, T., Rack, F. R., Saitzev, V., Sokolov, V., Makarov, A., Schwarze, S., Wunderlich, T., Wurr, K., and Boetius, A.: MOSAiC Extended Acknowledgement, Zenodo, https://doi.org/10.5281/zenodo.5541624, 2021.

Radhakrishnan, K., Scott, K. A., and Clausi, D. A.: Sea Ice Concentration Estimation: Using Passive Microwave and SAR Data With a U-Net and Curriculum Learning, IEEE J. Sel. Top. Appl., 14, 5339–5351, https://doi.org/10.1109/JSTARS.2021.3076109, 2021.

Ren, Y., Li, X., Yang, X., and Xu, H.: Development of a Dual-Attention U-Net Model for Sea Ice and Open Water Classification on SAR Images, IEEE Geosci. Remote Sens. Lett., 19, 1–5, https://doi.org/10.1109/LGRS.2021.3058049, 2022.

Ressel, R., Frost, A., and Lehner, S.: A Neural Network-Based Classification for Sea Ice Types on X-Band SAR Images, IEEE J. Sel. Top. Appl., 8, 1–9, https://doi.org/10.1109/JSTARS.2015.2436993, 2015.

Ressel, R., Singha, S., Lehner, S., Rösel, A., and Spreen, G.: Investigation into Different Polarimetric Features for Sea Ice Classification Using X-Band Synthetic Aper-

ture Radar, IEEE J. Sel. Top. Appl., 9, 3131–3143, https://doi.org/10.1109/JSTARS.2016.2539501, 2016.

Ronneberger, O., Fischer, P., and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, arXiv [preprint], https://doi.org/10.48550/ARXIV.1505.04597, 2015.

Simonyan, K. and Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition, arXiv [preprint], https://doi.org/10.48550/arXiv.1409.1556, 2015.

Singha, S., Johansson, M., Hughes, N., Hvidegaard, S. M., and Skourup, H.: Arctic Sea Ice Characterization Using Spaceborne Fully Polarimetric L-, C-, and X-Band SAR With Validation by Airborne Measurements, IEEE T. Geosci. Remote, 56, 3715–3734, https://doi.org/10.1109/TGRS.2018.2809504, 2018.

Soh, L.-K. and Tsatsoulis, C.: Texture Analysis of SAR Sea Ice Imagery using Gray Level Co-occurrence Matrices, IEEE T. Geosci. Remote, 37, 780–795, https://doi.org/10.1109/36.752194, 1999.

Song, W., Li, M., Gao, W., Huang, D., Ma, Z., Liotta, A., and Perra, C.: Automatic Sea-Ice Classification of SAR Images Based on Spatial and Temporal Features Learning, IEEE T. Geosci. Remote, 59, 9887–9901, https://doi.org/10.1109/TGRS.2020.3049031, 2021.

Ullah, H., Khaleghian, S., Kromer, T., Eltoft, T., and Marinoni, A.: A Noise-Aware Deep Learning Model for Sea Ice Classification Based on Sentinel-1 Sar Imagery, in: 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, 816–819, https://doi.org/10.1109/IGARSS47720.2021.9553971, 2021.

Wang, Y.-R. and Li, X.-M.: Arctic sea ice cover data from spaceborne synthetic aperture radar by deep learning, Earth Syst. Sci. Data, 13, 2723–2742, https://doi.org/10.5194/essd-13-2723-2021, 2021.

Zhou, Z., Siddiquee, M. M. R., Tajbakhsh, N., and Liang, J.: Unet++: A Nested U-Net Architecture for Medical Image Segmentation, in: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, arXiv [preprint], https://doi.org/10.48550/arXiv.1807.10165, 2018.

Zhou, Z., Siddiquee, M. M. R., Tajbakhsh, N., and Liang, J.: UNet++: Redesigning Skip Connections to Exploit Multiscale Features in Image Segmentation, IEEE T. Med. Imaging, 39, 1856–1867, https://doi.org/10.1109/TMI.2019.2959609, 2020.