

Estimating Sea Ice Concentration from SAR: Training Convolutional Neural Networks with Passive Microwave Data

Colin LV Cooke, *Student Member, IEEE*, and K. Andrea Scott, *Member, IEEE*,

Abstract—Historically, sea ice concentration has been measured through use of passive microwave sensors, as well as human interpretation of synthetic-aperture radar (SAR). Although passive microwave data is processed automatically, it suffers from poor spatial resolution and the higher frequency channels are sensitive to weather conditions. Deep learning has demonstrated its ability to perform complex, accurate, analysis of images; here, we apply deep learning to estimate ice concentration from SAR scenes. We developed a deep convolutional neural network (CNN) that predicts sea ice concentration from SAR, trained upon passive microwave data. The model achieves a 5.24%/7.87% error on its train and test set respectively. To assess the real-world applicability, we performed an independent validation on 18 SAR scenes (from two distinct geographical regions), not previously seen during training or test. Comparing against human generated ice analysis charts, we achieved an L_1 error of 0.2059, competitive with passive microwave ($E_{L_1} = 0.1863$) for the Canadian Arctic Archipelago. For the Gulf of Saint Lawrence region, we achieved an L_1 error of 0.2653, significantly better than the passive microwave result ($E_{L_1} = 0.3593$). By using novel techniques for model training, as well as training entirely upon passive microwave data, we present an accessible and robust method of developing similar systems for processing SAR¹. Our results suggest that with further post processing, CNNs are accurate and robust enough to be used for operational tasks.

Index Terms—Sea Ice, Neural Networks, Synthetic Aperture Radar, Remote Sensing, Image Processing

I. INTRODUCTION

Sea ice concentration is the proportion of a given area in the ocean that is covered by sea ice. Satellite sensors are the key instruments through which observations are made of sea ice concentration in the Arctic. These sensors differ in terms of their imaging properties, which impacts the spatial resolution of the data as well as the sensitivity of the data to atmospheric and surface conditions. The calculation of ice concentration from satellite imagery uses knowledge of physical properties of the observed signal, such as brightness and temperature from passive microwave data [1] or albedo from an optical sensor [2]. The observed signal can then be used to deduce if a given pixel is likely ice-covered, open water, or a mixture of the two, which would correspond to an intermediate ice concentration.

The most widely used sea ice concentration estimates are those from passive microwave radiometers, due to their relative

insensitivity to atmospheric moisture (at low frequencies), and solar illumination. However, a drawback of these observations is the trade-off between the spatial resolution and atmospheric contamination. The smallest available instrument field of view (IFOV) on a passive microwave sensor that is typically used for sea ice imaging is 3×5 km for the 89GHz channels on the AMSR2 sensor. At this frequency the satellite measurement is sensitive to atmospheric water vapour, cloud liquid water, wind-roughening of open water regions, and variable sea ice emissivity/temperature leading to difficulties accurately estimating ice concentration in the marginal ice zone (MIZ). [1, 3, 4]. As an alternative to passive microwave imagery, data from active microwave sensors, such as synthetic aperture radars (SARs) can be used to provide information of the ice conditions during all weather conditions. The spatial resolution of these images is much higher (approx 50m) than that of passive microwave data. However, it is more challenging to use SAR imagery for the task of ice concentration estimation than passive microwave imagery. This is due to the sensitivity of the signal to the imaging geometry and small-scale details of the surface conditions (e.g., variations in moisture, salinity and roughness). Typically, SAR images are manually interpreted at operational ice services by image analysis experts. This provides information on ice type and ice concentration over regions of the image, or polygons, that the analysts consider as having relatively homogeneous conditions through visual inspection of the imagery.

Preparation of image analysis charts is time-consuming, and it is highly desirable to automate this process. This issue is one of concern for operational ice centres due to increasing volumes of SAR data from the Sentinel missions, and anticipated data volumes from the RADARSAT constellation mission (RCM). Even at the present time, the large volume of SAR imagery precludes the preparation of an image analyses chart for every SAR image. The images are usually only generated for specific regions and seasons where there is known ship traffic. This may not seem like a significant limitation, but it means that the data are underutilised. For example, SAR sea ice imagery is not routinely used in operational data assimilation systems, unlike sea ice concentration from passive microwave sensors, which is widely used.

The task of estimating sea ice concentration from SAR imagery is an excellent example of one that can benefit from recent developments in deep learning, in particular the use of deep convolutional neural networks (CNNs). A CNN makes use of the convolution operation to learn features from images

Colin LV Cooke is corresponding author: clvcooke@edu.uwaterloo.ca

K. Andrea Scott is in the Department of Systems Design Engineering, University of Waterloo, Canada.

¹Code available: github.com/clvcooke/Estimating-SIC-from-SAR

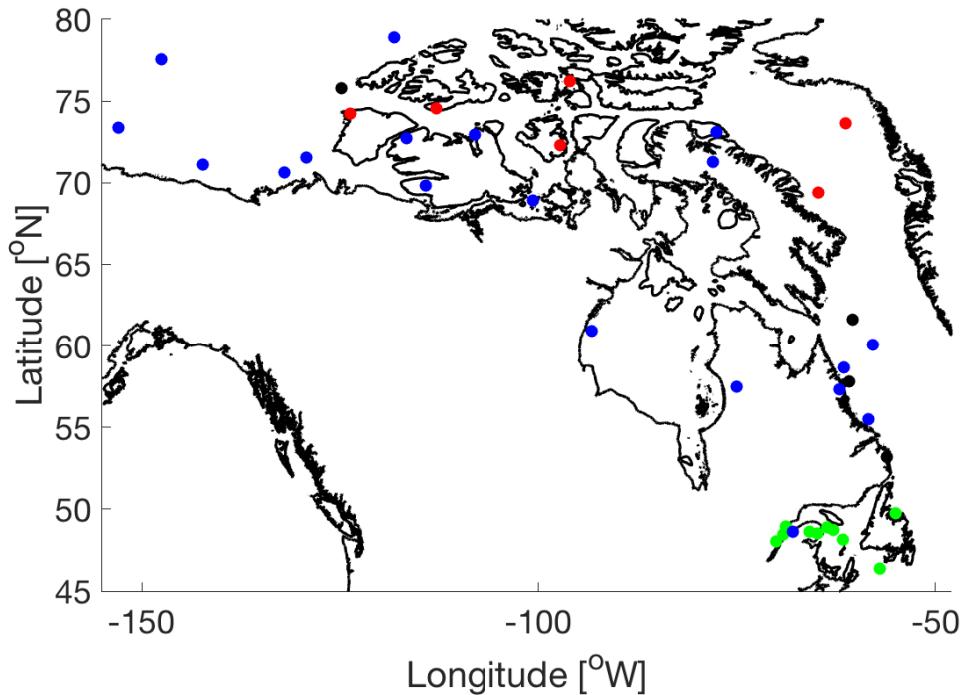


Figure 1. Geographic locations of SAR images used in the study. Blue symbols are images used in training, black symbols are test images, red symbols are evaluation images in the Canadian Arctic Archipelago, green symbols are evaluation images in the Gulf of Saint Lawrence.

and uses these features in a classification or regression task. A typical CNN consists of sequential layers of convolutional filters. It is common to place down-sampling operators after each convolutional layer to reduce the amount of information in the network. The outputs from this series of filtering operations are weighted to evaluate their relative importance with regards to the regression or classification task. This is achieved through one or several fully connected layers, typically with a rectified linear activation. The output of this weighting is the *classification* or *regression* result, either a *probability vector* or a scalar output. A broad review of deep neural networks as well as their applications and performance can be found in [5].

The method presented in this paper follows and extends two proof of concept studies [6, 7] that successfully estimated sea ice concentration from SAR using a CNN with ice concentration from image analysis charts as training data. This type of training data is very common in methods for supervised classification or regression from SAR sea ice imagery [8, 9, 10, 11, 12, 13, 14, 15, 16, 17]. However, the training data, image analysis charts, may not be well-suited to task of learning ice concentration. These charts contain a single ice concentration label for a region considered by an ice analyst as containing homogeneous ice conditions. The data are categorical, with labels in increments of 10% [10] or 20% [18]. When using these data in regression, it is commonly assumed that they are continuous valued [10, 12]. Another challenge with the use of image analysis charts is the fact that they are not available over a wide range of seasons or regions, as mentioned earlier.

For these reasons, the present study investigates the use

of passive microwave sea ice concentration to provide the labels for the neural network approach. This type of data is readily available on a daily basis over the entire Arctic region. However, one challenge with using this sea ice concentration is that passive microwave data is noisier than image analysis charts. This, in combination, with the noise present in SAR images, can lead to poor results. In the present study this is addressed using a novel approach to train the CNN. The proposed method is shown to be robust to noise through the use of a two-step training procedure that can be considered an example of *curriculum learning* [19] in the sense that the CNN is given a slightly easier problem to solve initially, followed by a more difficult one. It is also similar to the approach used in [20], where a two-step procedure was used to classify remote sensing images of urban areas.

The specific research questions addressed in this study are as follows:

- 1) Can we obtain a reasonable estimate of sea ice concentration from dual-polarization SAR imagery using a CNN with passive microwave sea ice concentration as training data?
- 2) How well does our method work when applied to images that contain ice conditions outside of the train/test data set?
- 3) What is the added value of ice concentration from SAR as compared to that from passive microwave data?

We evaluate the sea ice concentration (SIC) estimates from our CNN on a set of SAR images covering a wide range of geographic locations and seasons, including the challenging conditions of freeze up, and demonstrate SIC errors on two

independent evaluation data-sets that are comparable to, or better than, those from other studies. The data used in the study are described in Section II, methodology and experimental set-up are given in Sections III and IV and results and conclusions are in Sections V and VI respectively.

II. DATA

A. SAR imagery

The image dataset used for training and testing consists of 24 RADARSAT2 SAR scenes (Table I). An additional 18 RADARSAT2 SAR scenes were used for evaluation of the CNN SIC (Table III). Each scene consists of an HH (horizontal transmit polarization, horizontal receive polarization), and an HV (horizontal transmit polarization, vertical receive polarization) image. The images were obtained at an imaging frequency of 5.404 GHz over incidence angles ranging from 20 to 49 degrees in ScanSAR wide mode. The swath width is approximately 500km with nominal pixel spacing of 50m. These images are of the type currently used by operational ice services to produce ice charts for navigation. The available SAR images were captured from the regions shown in Figure 1. Image acquisition dates, central latitude and longitude, and general ice conditions for each image are given in Table I. Note the SAR images were obtained over a wide range of ice conditions, covering regions of first-year and multi-year ice, with some images during the freeze-up and melt seasons. This lends itself to a challenging problem, because both active and passive microwave sea ice signatures (e.g., typical values of brightness temperature or backscatter for a given ice condition) overlap with open water signatures over such a wide range of ice conditions [1, 21, 22]. The SAR images were geo-referenced to the earth's surface using code specifically written for this purpose following the documentation provided in [23]. Selected geo-referenced images were compared with those geo-referenced using the European Space Agency toolbox, SNAP, and were found to be equivalent.

B. Passive microwave sea ice concentration data

Sea ice concentration estimated using the Artist Sea ice (ASI) algorithm [24] was used to provide the labels for the CNN. The ASI algorithm fits the polarisation difference of the 89GHz channels to ice and water tie points assuming a cubic relationship. In the ASI algorithm, weather filters based on the lower frequency channels of the AMSR-E sensor are applied to remove spurious ice over open water that can result due to atmospheric water vapour or cloud moisture, which have spectral signatures close to those of sea ice. The objective of the first filter is to remove spurious ice concentration values due to cloud water. This is done by setting the ice concentration to zero for all points for which the 36.5/18 GHz gradient ratio (vertical polarization) is greater than 0.045. The objective of the second weather filter is to remove spurious ice concentration values due to water vapour, which is done by setting points for which the 23.8/18.7 gradient ratio is greater than 0.04. This weather filters has been found to remove ice concentrations less than 15% [24]. This means, these weather filters often lead to significant reduction of sea

ice concentration in the marginal ice zone, as can be seen when results are compared visually with other data sources, such as SAR imagery, optical data or ship observations [25, 2, 26]. The ASI SIC has errors that are comparable to other passive microwave retrieval algorithms, and was chosen because of its high spatial resolution (6.25km when brightness temperatures from AMSR-E are used, and 3.25km when brightness temperatures from AMSR2 are used) in comparison to other passive microwave SIC products (typical spatial resolution 10-25km). Since we are interested in navigable waters around Canada, such as the Labrador Coast and Canadian Arctic Archipelago, this high spatial resolution was critical to resolve the regions in narrow channels and bays.

C. Image Analysis Charts

For evaluation of the CNN output, SIC from operational image analysis charts is used. To prepare these charts, ice analysts draw polygons indicating areas where the ice conditions appear to be homogeneous. An ice concentration label is then assigned to each homogeneous region given in increments of 0.1. The preparation of image analyses is subjective [27], with errors in the position of the ice edge and polygon boundaries (± 500 m) as well as possible errors due to the interpretation of the SAR imagery. However, image analysis charts are considered an accurate source of information pertaining to sea ice conditions corresponding to SAR imagery [28]. For the present study rasterized image analysis charts were used, with a pixel spacing of ≈ 5 km.

III. METHODOLOGY

A. Working with Large Scale Information

The ScanSAR wide images used consist of approximately $10,000 \times 10,000$ pixels, over an order of magnitude greater than the images that typical CNNs are designed to use. Thus we cannot process an image in its entirety without down-sampling, a process that would obliterate many important details. In addition, a single label or regression result for the scene as a whole is of no interest. The images contain measurements of radar back-scatter, with each measurement corresponding precisely to a spatial position on the surface of the earth. It is the mapping from the back-scatter at a spatial location to a geophysical quantity, (e.g., ice concentration), that is of interest. To produce a map of this geophysical quantity, a process is required that enables processing of entire images, without significant decrease in spatial resolution or reduction in quality of results. Due to the speckle noise present in SAR images, both back-scatter and spatial context are needed to provide information of what is being imaged. Hence, pixelwise processing of the image information is not optimal.

Following [6], we utilised a patchwise processing method, where overlapping, square sections of the scene are extracted and processed individually. After processing, the results were re-combined and joined into a regression result for the entire scene. This method is similar to two dimensional convolution, but instead of performing matrix multiplication, each image patch is evaluated with the trained deep convolutional neural network (CNN) to estimate the ice concentration. Following

Table I

RADARSAT-2 IMAGE DATA USED TO TRAIN AND TEST THE MODEL. FYI DENOTES FIRST-YEAR ICE, MYI DENOTES MULTIYEAR ICE, MIZ DENOTES MARGINAL ICE ZONE AND OW DENOTES OPEN WATER.

	Acquisition date	Central coordinate	Ice conditions
Training	January 31, 2010	48.65°N, -67.90°W	Thin FYI and OW, freeze-up
	April 3, 2010	70.64°N, -132.10°W	Primarily FYI, some MYI
	April 5, 2010	71.10°N, -142.27°W	Primarily FYI, some MYI
	October 7, 2010	72.73°N, -116.70°W	FYI and MYI
	October 9, 2010	71.29°N, -78.06°W	MYI
	October 13, 2010	68.94°N, -100.73°W	Primarily OW
	December 5, 2010	60.92°N, -93.27°W	FYI with some OW
	December 13, 2010	73.09°N, -77.63°W	MYI
	February 13, 2011	60.1°N, -57.9°W	MIZ and wind roughened OW
	February 16, 2011	57.29°N, -62.11°W	MIZ and OW
	February 20, 2011	55.47°N, -58.53°W	MIZ and OW
	March 19, 2011	57.8°N, -60.88°W	MIZ and OW
	July 2, 2011	69.85°N, 69.85°W	Primarily FYI
	July 5, 2011	57.48°N, -74.99°W	OW
	July 17, 2011	78.87°N, -118.2°W	MYI
	July 30, 2011	72.97°N, -108.09°W	MYI and OW
	September 3, 2011	77.56°N, -147.52°W	FYI, MYI and OW
	September 3, 2011	73.36°N, -152.86°W	FYI, MYI and OW
	September 3, 2011	75.78°N, -124.95°W	FYI, MYI and OW
	September 3, 2011	71.55°N, -129.28°W	MYI and OW
Testing	October 9, 2010	71.29°N, -78.06°W	OW
	February 17, 2011	53.18°N, -56.08°W	MIZ and OW
	February 23, 2011	48.79°N, -63.66°W	primarily OW
	April 4, 2011	61.62°N, -60.43°W	MIZ and OW

estimation we reconstructed the scene and performed spatially-aware post-processing. This process is outlined in Figure 2.

B. Assigning labels to the SAR image backscatter

For a typical image classification problem a single label for each input image is given (e.g., dog, tree, car). In contrast, for the present application, the labels, which correspond to ice concentration, are continuous-valued data, with a range of values over the image space. Hence we must both transform the learning problem from classification to regression (discussed in Section III-D) and we must transform the map of sea ice concentration data to fit the given SAR scene. For maximum flexibility we construct an *ice concentration* representation of the SAR scene, the same aspect ratio as the input SAR scene (although at a lower spatial resolution). Thus, when we look to see the concentration of a point within the scene, we can reference this representation directly in grid space.

To construct this representation, we iterate over equally spaced pixels in the SAR scene and calculate the corresponding ice concentration at the pixel using Algorithm 1, an inverse weighting of the ice concentration values of the four nearest neighbours in the corresponding passive microwave image. This provides a label for each pixel of the SAR image and theoretically allows for the resolution of the output to be the same as that of the input.

Algorithm 1: Interpolation Scheme

Data: SAR data-point, PMW data-points
Result: Concentration value for SAR data-point
 $LatLon_{SAR} = getLatLon(data_{SAR})$
 $LatLon_{PMW} = getLatLon(data_{PMW})$
 $dists_{PMW} = norm(LatLon_{SAR} - LatLon_{PMW})$
 $indices_{PMW} = argsort(dists_{PMW})$
 $concentration = 0$
 $totalDistance = 0$
for each i in $[0, 1, 2, 3]$ **do**
 $concentration += \frac{data_{PMW}[indices_{PMW}[i]]}{dists_{PMW}[indices_{PMW}[i]]}$
 $totalDistance += dists_{PMW}[indices_{PMW}[i]]$
end
 $concentration = \frac{concentration}{totalDistance}$
return concentration

C. Masking of land pixels

Note that because our focus is on images through navigable Canadian waters, many of the images used to both train and test the CNN model had significant land cover. To work with this land and still produce accurate results a land mask was used to identify image patches containing land. Patches that were entirely, or centred, on land were excluded from the dataset. For the remaining patches, a value of one was used to mask over all land within an image patch. In this way the model was trained with a consistent representation of land.

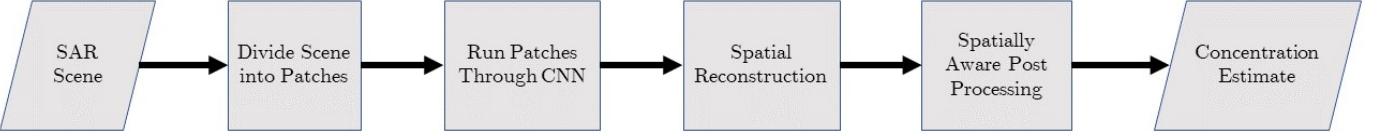


Figure 2. Image processing pipeline: Original SAR scene is broken up into patches and run through the deep convolutional neural network (CNN). Following this processing, the scene is reconstructed to form an image of the same aspect ratio, but with a smaller overall size, containing the concentration information. This image is post-processed with a spatially-aware post-processing scheme to provide the final result.

D. Convolutional Neural Network Usage

The CNN chosen for use in the present task is *DenseNet* [29]. This CNN performs extremely well on the challenging ImageNet database, which is a database consisting of millions of images, primarily of everyday objects and entities. *DenseNet* achieves a similar classification accuracy as other leading CNNs, but uses less than half the number of parameters. This is due to the unique architecture of *DenseNet*. In a typical sequential convolutional neural network (where there is a single pathway for information) multiscale information is passed sequentially through layers of convolutional filters, rectified linear activations and down-sampling operations [30, 10, 31]. In *DenseNet*, skip connections (direct connections that skip over sequential blocks of layers) are used to directly pass the information from each layer to a concatenation layer, allowing the network to use features of multiple scales easily. These connections are an especially important attribute in remote sensing where important features are seen at all scales. For sea ice SAR imagery, this range of scales spans the fine grained texture of open water, to larger, linear features (that may be ridges) or openings (e.g., leads) within the consolidated ice cover. The simultaneous use of information at all scales is also similar to the manner in which ice operators carry out the manual analysis of SAR sea ice imagery, and more generally to the manner in which humans interpret SAR images [32]. Another advantage of *DenseNet* in comparison to other CNN architectures, in particular for remote sensing applications where training data sets are limited in size, is that the direct skip connections allow more efficient use of the training data, since the gradients computed during back-propagation are more accurate and less noisy.

Although *DenseNet* was designed for image classification, it is easily modified to perform regression. The final softmax classification layer in *DenseNet* is removed and replaced with a single non-linear activation (*sigmoid*). The three channel input (typically RGB) is replaced with a two channel image, consisting of HH and HV channels. An input shape of $221 \times 221 \times 2$ was initially used, with further tests carried out using $321 \times 321 \times 2$. When training this network, we seek to minimise absolute error, which is the absolute value of the difference between the target label and the model output.

IV. EXPERIMENTAL DESIGN

A. Chosen Patch Size

No filtering or averaging was done over the SAR images; hence each patch contributes to a single 1×1 output, centred within its corresponding patch. When there is no overlap

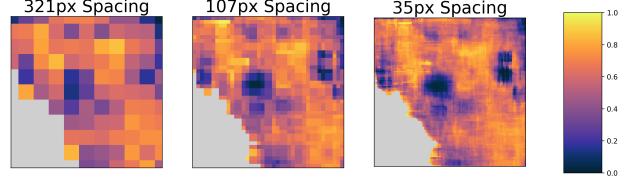


Figure 3. Sample output results for the various patch spacings tested: Note that as the patch spacing decreases the details of the ice cover become finer, but overall the impact on the results is not significant. This indicates smaller or larger spacing could be chosen based upon the application.

between the patches, this yields ice concentration with a spacing of $11 \text{ km} \times 11 \text{ km}$ for the patches of size 221×221 , and a spacing of $16 \text{ km} \times 16 \text{ km}$ for the patches of size 321×321 . To achieve a higher resolution output we used a spacing between patch centers of 73 pixels and 64 pixels respectively for the 221×221 and 321×321 patch sizes. This corresponds to an output at intervals of 3.68 km for the 221×221 patch, and 3.21 km for the 321×321 . This is similar to the spacing of recent 89GHz ice concentration estimates from AMSR2 (3.25km), and finer than the AMSR-E ice concentration used here for training data (6.25km) [24] to provide the labels. We observed that the accuracy of the output was insensitive to patch spacing (Figure 3), suggesting a higher resolution output could be achieved, although we did not pursue this due to computational constraints.

B. Test/Train Split

We split the data into a training set and a test set using roughly a 4:1 ratio, resulting in 19000 training samples and 5000 test samples. To allow for truly independent samples in the test set we took a subset of the full SAR scenes and used them to develop the test set, see Table 1. These scenes were not included in the training set. The test set scenes were drawn randomly from the available set, however the scenes were validated to represent a variety of ice and water conditions. Although splitting the data into train/test in this manner is useful for gauging generalisation during training, we simultaneously lose some perspective into the progress of the model. This is due to inherent within-image bias that leads to over or under representation of features for a given scene. Here, this issue was minimised by choosing a test set that contained new, thin ice and thicker multiyear ice, in addition to calm and rough water. Table I outlines the general distribution of ice conditions throughout our chosen dataset as well as the conditions represented within the test set.

C. Training Procedure

The training of the CNN model used in this study presented unique challenges. Upon training using the standard approach, where the image patches are fed directly into the CNN, it was found the model over exploited the correlation between the image texture and ice concentration present in some of the images, and falsely generalised this concept. This led to spurious ice over open water, and a noisy ice cover in regions of consolidated ice. This was a sign that the model may have been over-fitting and that either additional data or a suitable alternative to it was necessary.

As the DenseNet model had shown itself to be highly adept on challenging tasks prior; we chose to focus on the training to fix the over exploitation of texture. There are two typical methods of modifying the training to fix over-fitting: increase the size of the dataset, or augment the current data set [33, 34]. We chose data augmentation, and applied additive, zero centred, Gaussian noise to the input patches at training time.

In typical cases of data augmentation, the final model results directly from training on the augmented data, however in our case the augmentation process made the problem harder. While the results sometimes showed improvement from the previous iteration (no augmentation), the results were sub-optimal in certain circumstances. For example, in some cases the CNN ignored texture to the point of labelling smooth ice as open water (results will be shown in Section V). To improve upon these preliminary results, we applied a novel technique: after the augmented model had converged, we removed the augmentation (i.e., noise), reduced the learning rate and brought the model to a new minimum, forming a two-stage training procedure. This procedure allows the model to use the texture of the image whilst preventing over fitting upon it. It takes advantage of the regularisation provided by the additive noise in the early stages of training, which allows the model to learn with little risk of over exploitation of the texture, without removing the relevance of this concept entirely. The results of this procedure is described in Section V.

V. RESULTS

A. Training results

With our best model (*Prime-Large*) we were able to achieve an average absolute error of 5.25% *on the training set* and 7.87% *on the test set*, which is comparable to errors of 7-8% in other similar studies [12, 6]. To arrive at this result we underwent several iterations of training the model, in particular with regard to the amount of noise used and the size of the input patch. We will reference each iteration, including the iteration with no augmentation, with names shown in Table II. Select results for each iteration are displayed in Figure 4.

Table II
SUMMARY OF TRAINED MODELS WITH CONDITIONS THAT WERE USED IN
THE DEVELOPMENT OF THE PROPOSED METHOD.
*MODELS USED FOR EVALUATION

Name	Training Conditions	Patch Size
Base	No Augmentation	221×221
Base-Large	No Augmentation	321×321
Noise-A	Gaussian Noise ($\sigma = 0.15$)	221×221
Noise-B	Gaussian Noise ($\sigma = 0.10$)	221×221
Noise-Large	Gaussian Noise ($\sigma = 0.10$)	321×321
Prime-A	Initialised with <i>Noise-A</i>	221×221
Prime-B*	Initialised with <i>Noise-B</i>	221×221
Prime-Large*	Initialised with <i>Noise-Large</i>	321×321

1) *Base Models*: We trained two base models using recommended procedures of early stopping [35], coupled with learning rate reduction [36]. These models obtained respectable results, however showed signs of failing to generalise and possibly over-exploiting aspects of the training set. This led to noisy ice cover in some cases, and complete misinterpretation of ice as water in others. The move from a patch size of 221 × 221 to 321 × 321 showed significant improvement but still did not yield satisfactory results. Figure 4 row 3 (Base) highlights a strong water prediction but a significant low bias. Figure 4 row 4 (Base-Large) fixes the low bias but still yields patchy results within the largely consolidated ice.

2) *Noise Models*: We trained three models, with the same training set as the Base models however we augmented the images with Gaussian noise. The Gaussian noise was applied in a *dynamic* manner, meaning it is drawn from a distribution at the training time such that different noise is applied to each image patch, and different noise is applied each time the patch was processed (e.g., each epoch). This was done to prevent the model from memorising or over-fitting to irrelevant details. We experimented with two noise levels, $\sigma = 0.15$ and $\sigma = 0.10$ (image channels were normalised to have pixel values between 0 and 1). The higher noise level led to a complete irrelevance of texture information within the images in the context of ice concentration estimation. The regression often yields a value of 0 (indicating water) over consolidated ice, while the lower noise level showed partial salience (row 5 vs. row 6 in Figure 4). When we combined the lower noise level ($\sigma = 0.10$) with a larger patch size we saw a further improvement (rows 6 vs. row 7 in Figure 4). Shown in row's 6 and 7 we also see that the model lacks the ability to predict medium concentrations well, often outputting values at the extremes (values of zero and one).

3) *Prime Models*: Following each of the three noise models we developed a pretrained noiseless model. This was done by using the weights of the relevant model trained with noise to initialise a new network. This new network was then re-trained without noise injection and a reduced learning rate. Throughout all models we see partial to full recovery of texture relevance. For *Prime-A* there are still many instances of inaccurate regressions for large swaths of the scenes. For *Prime-B* we see mostly accurate regressions, however the ice cover is noisy and contains small "holes" of ice misclassified

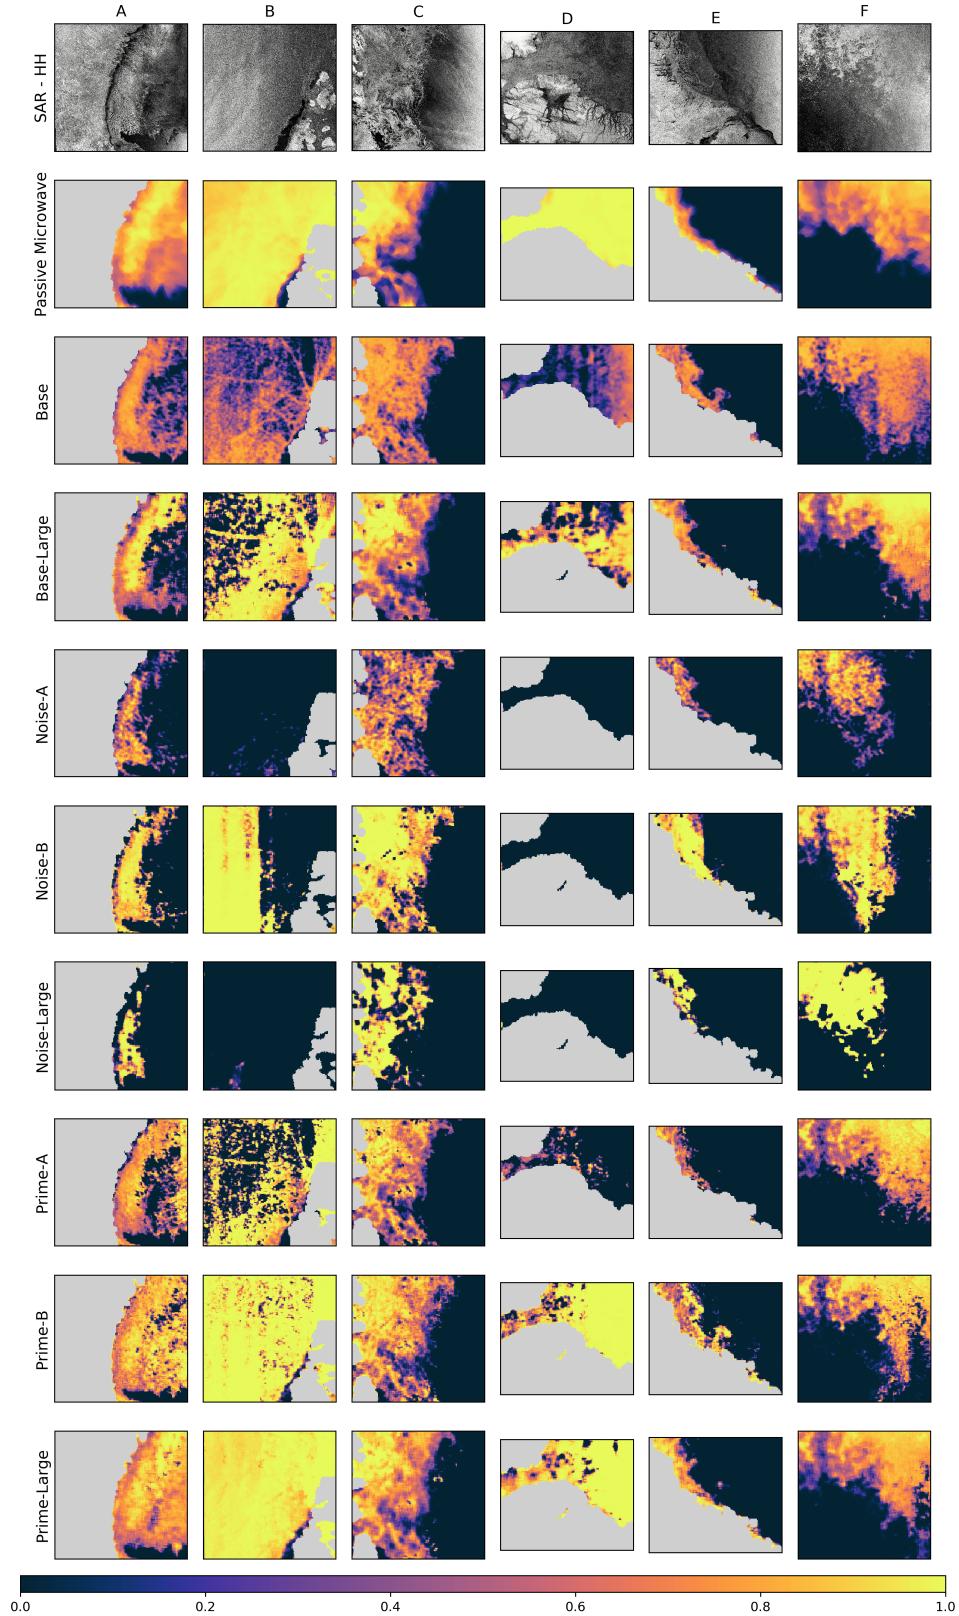


Figure 4. Select results of model iterations. Similar results are found for other training images. Moving from left to right for each column we see first year ice with some open water (A), majority multi-year ice (B), primarily first year ice with some multi-year ice (C), majority multi-year ice (D), a marginal ice zone as well as open water(E), and a mix of first year ice, multi-year ice and open water (F). Calling attention to column E we can see that two of the models (Base and Prime-B) detect an ice eddy absent off the coast. This ice eddy is absent in the passive microwave version of the scene, and not represented in the training data. However, the eddy can be seen in both the SAR imagery and the manual image analysis chart.

as water. For the final model, *Prime-Large*, we combined the successes of *Prime-B* with a large patch size. These final results showed significantly smoother ice cover combined with accurate regressions over all inputs. We selected both *Prime-B* and *Prime-Large* for independent evaluation.

B. Spatially Aware Post Processing

As shown in the *Prime-B* and *Prime-Large* rows of Figure 4, in some cases the the ice cover for these results is noisy, with holes over large regions of what is likely consolidated ice. These holes may arise because the method does not take into account spatial context; without information about the space surrounding each input patch, the model has no way to relate information about adjacent patches. This could be more obvious in the present study than Wang et al. [6, 7] because the training data, ASI ice concentration, is much noisier than the ice analysis charts used in previous studies.

To obtain a smoother ice cover post-processing was applied to the ice concentration estimated by the CNN. This consisted of a median filter acting over a region of 5×5 pixels (excluding sections of the result that were covered by the land mask). Examples of the post processed images are shown in Figure 5. It can be seen that the method reduces the incidence of holes in the ice cover while maintaining a crisp ice edge and excellent features. For example, the ice eddy that can be clearly seen in the SAR imagery off the coast of Labrador (and is identified in the manual image analysis chart) is retained with post processing (see column 5 in Fig. 5).

C. Model Evaluation

For the proposed method to be useful for a wide variety of applications, the model must generalise to a wide variety of ice conditions, some of which may not be represented well within the training set. To evaluate this capability of our method, we ran the CNN in test mode over two additional sets of SAR sea ice imagery that are distinct from the images used in the training or test procedure. We refer to these as our evaluation data sets. The first set of images were acquired in the Canadian Arctic Archipelago during September and November of 2013, while the second are a set of SAR images acquired during freeze-up in the Gulf of Saint Lawrence, 2014. Locations of the image acquisitions are shown in Figure 1.

For both sets of test results, the ice concentration from the CNN was compared with that from image analysis charts provided by the Canadian Ice Service. The statistics used for comparison were the mean (E_{bias}), mean absolute difference, (E_{L1}) and standard deviation of the difference, (E_{std}), where the difference is between the ice concentration from image analyses and the ice concentration of either the CNN or ASI dataset. The sums used in the statistics were taken by generating a single vector of differences for each dataset (i.e. concatenating the differences for each day). Note that some of the differences between the CNN or ASI ice concentration and that from the image analyses is due to the fact that the image analyses have homogeneous ice regions within polygons, with abrupt changes at polygon boundaries. These abrupt changes can accentuate measures that are sensitive to outliers, such as the E_{std} .

1) *Arctic Images 2013*: The CNN model was evaluated using six SAR scenes acquired in the Canadian Arctic in September-November, 2013. Overall scores for all six scenes are shown in Table IV. It can be seen that the post processing has little impact on E_{bias} but decreases E_{L1} and E_{std} . For the data-sets with post processing the overall differences are similar to those from the ASI data, with lower E_{bias} and larger E_{std} and E_{L1} . However, visualisation of the results shows some interesting differences, as can be seen in Figures 6, 7 and 8, representing good, medium and slightly poorer CNN results.

For the first set of results (Figure 6), it can be seen that the CNN ice concentration is very similar to that from the ASI, and both are close to the image analysis chart, although all methods miss the region of lower ice concentration close to the north-west corner of Banks Island (red circle in Figure 6a). A close inspection of the SAR image showed this region is sparsely covered by isolated floes, which were not present in the training images. For the second set of images (Figure 7), it can be seen the CNN tends to misrepresent some of the dark tones in the SAR image as regions of low ice concentration, an aspect that is improved with post-processing. The CNN does a good job capturing the ice cover close to the ice edge (green circles in Figure 7a). Again, all methods fail to capture the thin new ice (red circle in Figure 7a). For the third set of images, shown in Figure 8, it can be seen the CNN does an excellent job capturing the ice on the north side of Prince of Whales Island, but misses a great deal of the low concentration ice, in particular the isolated floes in the top right of the image, and some of the diffuse ice cover the the bottom right (latter indicated by the red circles in Figure 8). Note that CNN output is not sensitive to the banding in the HV image, which is different from previous results using a similar method [7].

2) *Gulf of Saint Lawrence, freeze-up 2014*: The Gulf of Saint Lawrence (GSL) dataset is particularly challenging because the ice conditions represent freeze-up, with large regions of thin new ice. It is also a challenging regime for passive microwave retrieval algorithms, which typically underestimate the ice concentration for thin ice due to the fact that the emissivity of thin ice at high ice concentration is similar to that corresponding to thicker ice of an intermediate ice concentration. Overall scores for the GSL dataset are given in Table V. It is very impressive that the CNN is able to improve upon the ASI SIC, having both a lower standard deviation and lower bias of differences with the image analyses both with and without post processing. Note that these errors are comparable to, or lower than, those in [37], where ice/water observations from the SAR data were combined with ASI SIC using a data assimilation approach.

Some sample results are shown in Figures 9 and 10. In Figure 9 it can be seen that ASI significantly underestimates the ice concentration over a larger portion of the domain than the CNN. On this date the ice cover is primarily new ice and thin ice, but the dynamic features in the ice cover (such as ice eddies and ice filaments) allow the CNN to do a reasonable job estimating ice concentration, in particular the ice edge (green circle in Figure 9a). In Figure 10 both methods show generally good agreement with the image analysis, with

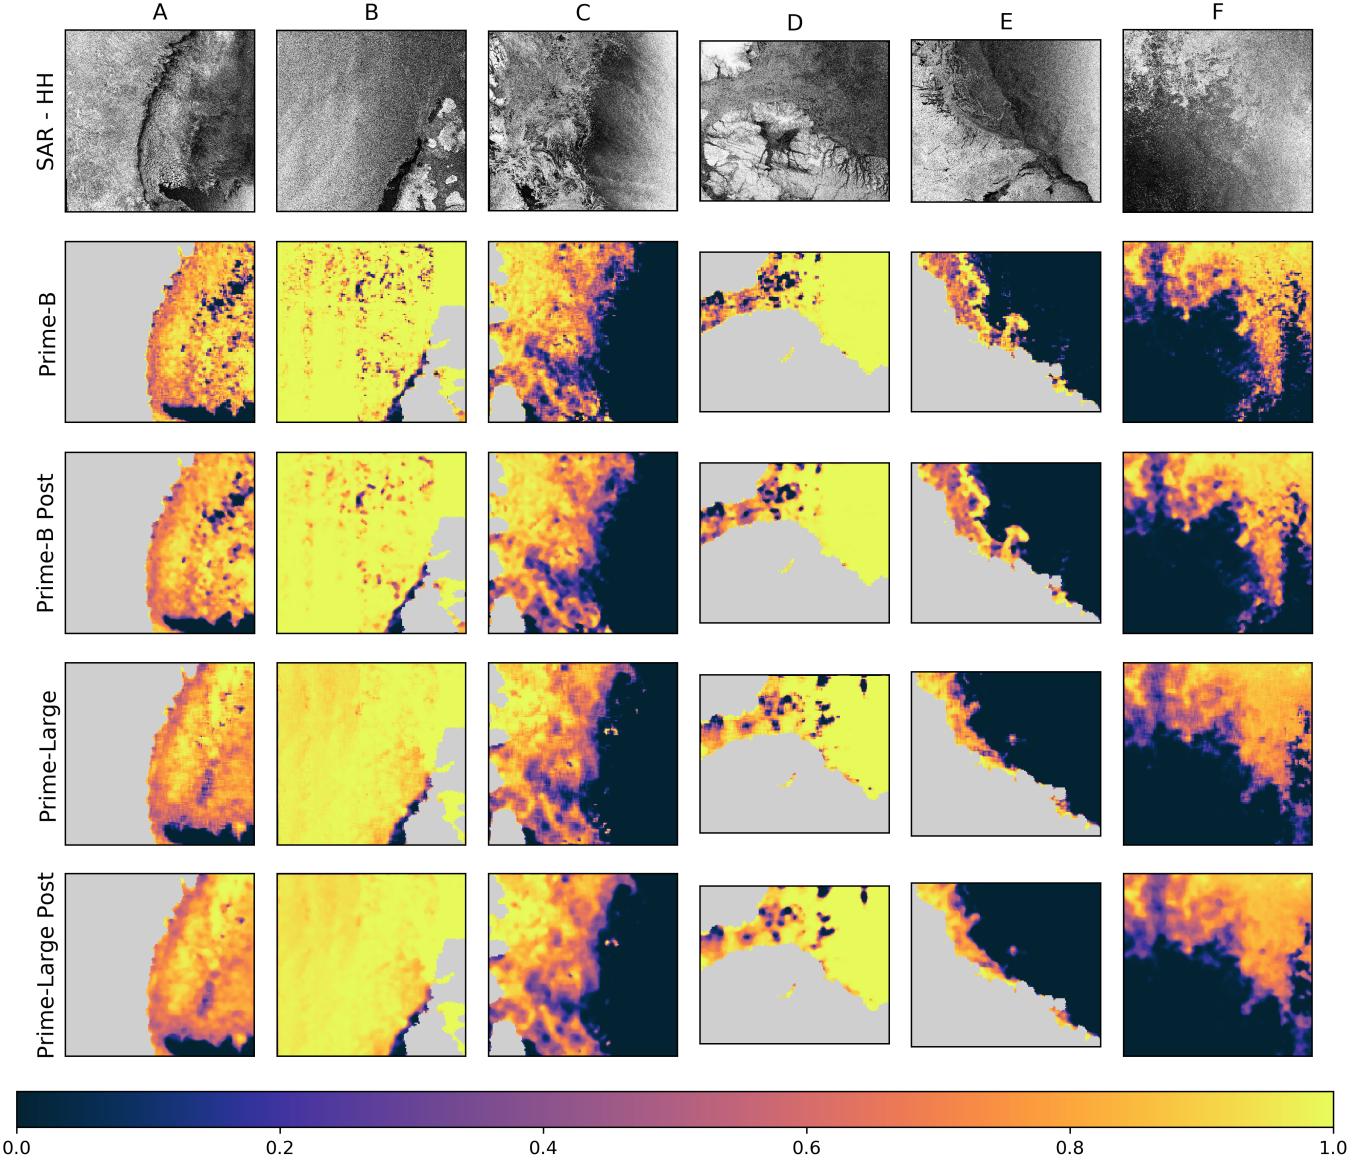


Figure 5. Impact of spatially aware post processing on the results: It can be seen that although the post processing reduces the noise in the results it does not radically change them. While this denoising is visually attractive and allows the resultant representation to be more easily interpreted it does not substantially change the accuracy of the result for any important metric.

the CNN producing better estimates north-east of Anticosti Island (green circle in Figure 10a). Both underestimate the ice concentration for the thin, new, ice that is in the region northwest of Prince Edward Island (red circle in Figure 10a), which appears as dark in both the HH and HV images with little texture.

3) *Summary of Model Evaluation:* Overall these results demonstrate the added value of ice concentration from SAR imagery as compared to that from passive microwave data: the CNN output contains a better representation of the ice edge, and a clean representation of water without the use of a weather filter. Such a filter is sometimes necessary in the retrieval of information from SAR sea ice imagery due to the similarity of the appearance of wind-roughened open water and some sea ice conditions [22]. The holes in the ice cover

seen in some CNN results are reduced when a larger patch size is used. In addition, we note that features not well represented in ASI data used to provide labels to the CNN (out-of-sample data), such as isolated floes and in some cases, thin, new ice, are not as well represented by the CNN as other features.

VI. DISCUSSION

A. Advantages

The method presented, training a convolutional neural network (CNN) with passive microwave data, is able to predict SIC at a higher spatial resolution than the training data: a unique aspect of the design. As shown Figure 10 we are able to accurately show fine details of the ice edge without any examples of such edges within the training set. This

Table III

RADARSAT-2 IMAGE DATA USED TO EVALUATE THE MODEL. FYI DENOTES FIRST-YEAR ICE, MYI DENOTES MULTIYEAR ICE, MIZ DENOTES MARGINAL ICE ZONE AND OW DENOTES OPEN WATER.

	Acquisition date	Central coordinate	Ice conditions
GSL 2014	January 17, 2014	48.76°N, -65.39°W	NI and OW
	January 22, 2014	54.32 °N, -53.40°W	Thin FYI, MIZ and OW
	January 23, 2014	49.01°N, -67.36°W	Thin FYI and OW
	January 24, 2014	48.59°N, -64.51°W	NI, thin FYI and OW
	January 24, 2014	48.19°N, -61.59°W	NI, thin FYI and OW
	January 27, 2014	49.21°N, -66.96°W	NI, thin FYI and OW
	January 31, 2014	48.10°N, -63.76°W	NI, thin FYI
	January 31, 2014	48.26°N, -60.81°W	thin FYI and wind-roughened OW
	February 7, 2014	48.33°N, -62.89°W	NI, thin FYI
	February 8, 2014	49.22°N, -66.96°W	MIZ and wind roughened OW
	February 9, 2014	48.80°N, -68.22°W	thin FYI and OW
	February 10, 2014	47.66°N, -62.14°W	NI and thin FYI
Canadian Arctic 2013	September 10, 2013	72.89°N, -96.23°W	FYI, MYI, OW
	September 11, 2013	74.13°N, -112.16°W	FYI, MYI, OW
	September 13, 2013	72.20°N, -99.53°W	FYI, MYI, OW
	September 13, 2013	72.50°N, -126.12°W	OW and MYI
	November 12, 2013	64.80°N, -65.71°W	FYI, MYI and OW
	November 12, 2013	64.80°N, -65.71°W	FYI, MYI and OW

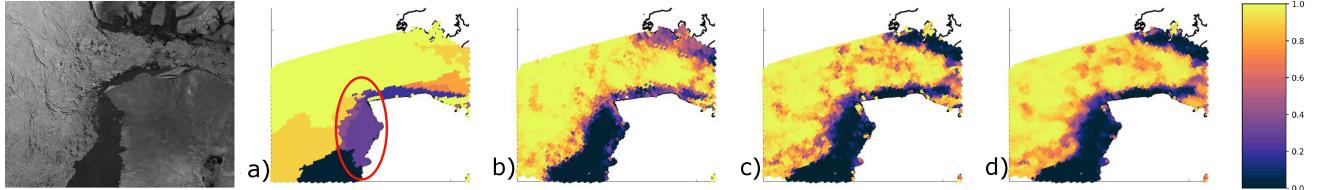


Figure 6. Results from scene acquired September 13th, 2013 in the Beaufort Sea near Banks Island. Central latitude and longitude of the image, 64.3 N and -139.7 W. Panel a) SIC from image analysis chart b) SIC from ASI, c) SIC from CNN - Prime B d) SIC from CNN - Prime B with post-processing. The HH SAR image is shown in the far left panel. The red circle indicates an ice detail not represented within the SIC from ASI or CNN.

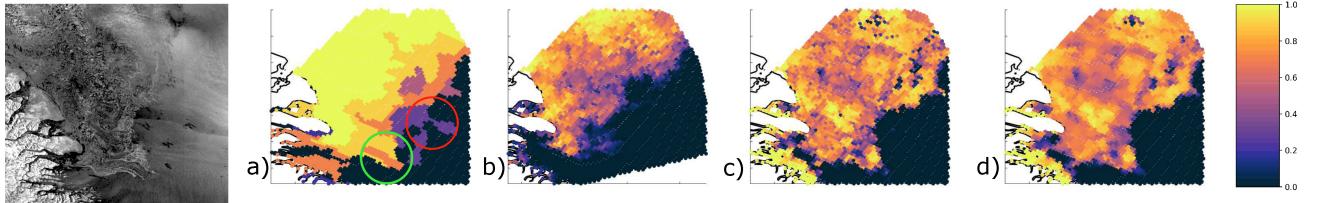


Figure 7. Results from scene acquired November 12th, 2013 in Baffin Bay. Central latitude and longitude of the image, 64.8 N and -65.7 W. Panel a) SIC from image analysis chart b) SIC from ASI, c) SIC from CNN - Prime B d) SIC from CNN - Prime B with post-processing. The HH SAR image is shown in the far left panel. The red circle indicates an ice detail not represented within the SIC from ASI or CNN. The green circle represents a detail replicated by the CNN and not the ASI.

is a testament to the flexibility of the *patchwise* method used. In our implementation of this method, the stride is a runtime parameter, which allows arbitrary spatial resolution on the output. For the present implementation, the stride was fixed over the domain, but in the future it could be made variable, which would allow for finer spatial resolution in regions of interest. In addition, for the images used in this study, we used no auxiliary information, such as wind-speed, to reduce spurious ice concentration retrievals over wind-roughened open water. This is important because wind-speed is very difficult to estimate accurately in the MIZ.

The method uses popular open source software, that is easy to configure for use with other SAR images, such as those from Sentinel-1; thus all data sources necessary to run the algorithm (passive microwave ice concentration, SAR imagery and DenseNet) are freely available. To encourage others to use this method, we have made the Python code used to generate the results in the paper available for download (URL will be provided in final manuscript). The method made use of GPU accelerated software, enabling the prediction of a 10,000x10,000 pixel SAR Scene in 160s by a mid-tier *NVidia Quadro P5000 GPU*.

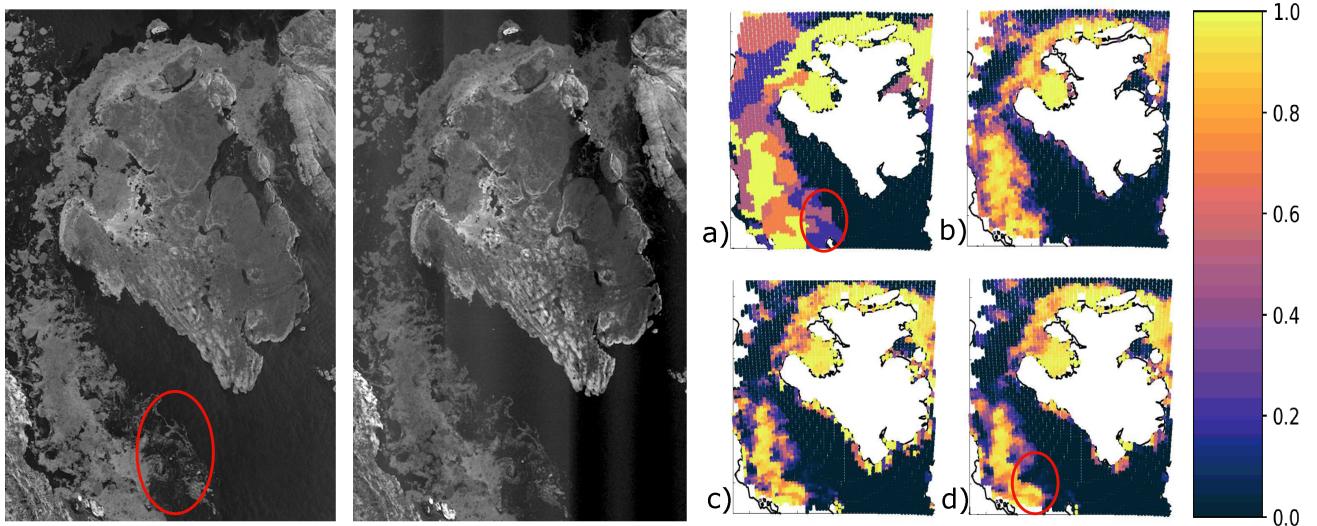


Figure 8. Results from scene acquired September 10th, 2013, near Prince of Wales Island, Canadian Arctic Archipelago. Central latitude and longitude of the image, 72.5 N and -102.5 W. Panel a) SIC from image analysis chart b) SIC from ASI, c) SIC from CNN - Prime B d) SIC from CNN - Prime B with post processing. The HH SAR image is shown in the far left panel. Note that the vertical stripes in the HV image are not reproduced by the CNN. The red circle indicates an ice detail not represented within the SIC from ASI or CNN.

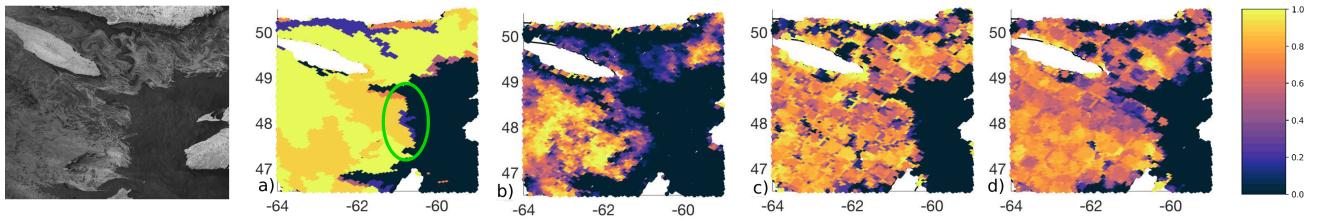


Figure 9. Ice concentration for SAR scene acquired on January 24, 2014 in the Gulf of Saint Lawrence. Panel a) SIC from image analysis chart b) SIC from ASI, c) SIC from CNN - Prime B d) SIC from CNN - Prime large. The HH SAR image is shown in the far left panel. The green circle represents an ice edge replicated well by the CNNs but not from the ASI.

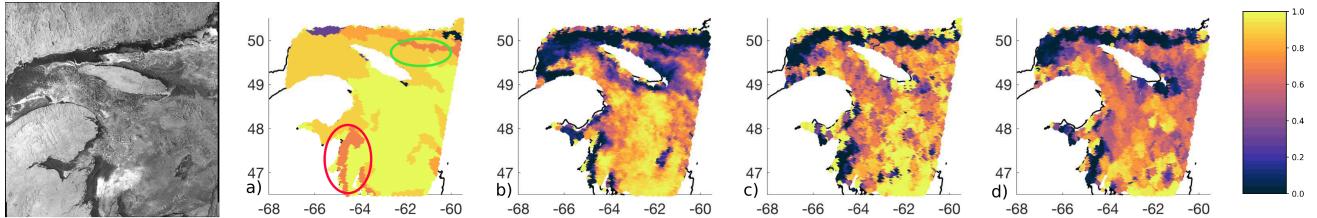


Figure 10. Ice concentration for SAR scene acquired on February 7, 2014 in the Gulf of Saint Lawrence. Panel a) SIC from image analysis chart b) SIC from ASI, c) SIC from CNN - Prime B d) SIC from CNN - Prime large. The HH SAR image is shown in the far left panel. The red circle indicates a region of thin, new ice that is not represented in the ASI or CNN SIC. The green circle indicates a region of relative agreement between the ASI and CNN representations, however the CNN is closer to the image analysis chart.

B. Potential Improvements

When predicting the concentration of sea ice, spatial context is important and provides information regarding how features should be interpreted. In this regard, sometimes a larger patch size leads to better incorporation of spatial context [7]. However, the drawback of a larger patch size is that it becomes increasingly difficult to have accurate predictions close to land. Instead of modifying the patch size, we propose a simple post-processing method, consisting of a median filter. A more sophisticated method would lead to better results, but we leave this to a future study, as the emphasis here is on

presenting the two-step training approach. For example, our results may benefit from historical data (e.g. climatology) to ground the predictions of the CNN model in realistic, regional expectations of ice patterns. Consider that a trained ice analyst may readily assign a non-zero ice concentration to a dark region downwind of an island based on their experience that it is probably new ice, whereas an automated analysis based upon a single patch has no such information.

We found the surprising result that although the larger patch size model (Prime-Large) performed better on the training set (when compared to Prime-B), it performed worse upon the evaluation set. This difference in performance may be

Table IV

TEST SCORES FOR IMAGES FROM THE CANADIAN ARCTIC ARCHIPELAGO, 2013. ERRORS ARE COMPUTED USING IMAGE ANALYSIS CHARTS AS VERIFICATION DATA. STATISTICS ARE CALCULATED USING 6 SAR SCENES ACQUIRED BETWEEN SEPTEMBER 10TH AND NOVEMBER 12TH, 2013. LOCATIONS OF THESE SCENES ARE SHOWN BY THE RED DOTS IN FIGURE 1

Dataset	E_{bias}	E_{L1}	E_{std}
Prime B	0.1033	0.2173	0.3377
Prime B (post)	0.1126	0.2059	0.3159
Prime large	0.1127	0.2327	0.3472
Prime large (post)	0.1153	0.2278	0.3385
ASI	0.1401	0.1863	0.2689

Table V

TEST SCORES FOR IMAGES FROM THE GULF OF SAINT LAWRENCE, 2014. THE SCENES CHOSEN WERE DURING THE FREEZE-UP PERIOD. ERRORS ARE COMPUTED USING IMAGE ANALYSIS CHARTS AS VERIFICATION DATA. STATISTICS ARE CALCULATED USING 12 SAR SCENES FROM THE PERIOD OF JANUARY 21ST TO FEBRUARY 10TH, 2014. LOCATIONS OF THESE SCENES ARE INDICATED BY THE GREEN DOTS IN FIGURE 1.

Dataset	E_{bias}	E_{L1}	E_{std}
Prime B	0.1147	0.2721	0.3858
Prime B (post)	0.1227	0.2653	0.3680
Prime large	0.1580	0.2983	0.3838
Prime large (post)	0.1660	0.2980	0.3773
ASI	0.3065	0.3593	0.3744

explained by the passive microwave data itself being worse than the Prime-B model upon the evaluation set; meaning if the Prime-Large model fits the passive microwave data more it also learns its flaws. We hypothesise that this difference could be remedied by fine-tuning the Prime-Large model with ice analysis charts as training data. Fine-tuning involves modifying the weights of the final layers in the model (the *concatenation layers*) to efficiently tune the model to a new dataset (in this case ice concentration charts). This would also help in the representation of new thin ice, which is represented fairly accurately on these charts. To emphasise the results that can be obtained without this method, we leave the examination of such fine-tuning to future studies.

In addition to possible post-processing, further computational resources could be committed to the prediction phase of the modelling, allowing for higher resolution outputs to be obtained. Although in some parts of the scene this may not hold any relevance, the method developed allows for variable resolution throughout the evaluation of the scene. This could be done iteratively, automatically refining the ice edge and other features of the evaluated scene.

VII. CONCLUSIONS

We have shown that sea ice concentration can reliably be estimated from SAR imagery using a CNN with only passive microwave data as training data. This is different than previous studies [6, 12] that have used ice charts as training data. We have found the use of ASI ice concentration can lead to a noisy or completely incorrect ice cover, and have developed a novel approach to reduce this noise. Dynamic, zero centred, Gaussian noise allowed aggressive pre-training with

little risk of over-fitting. The use of a *patchwise* evaluation method allows for variable output resolution, only restricted by the (meaningful) resolution of the input data. Finally, the CNN model developed in this study achieved competitive performance on ice conditions outside of its training and test conditions, providing an entirely independent evaluation of its performance. However, we did note that some specific ice features (e.g., isolated floes and thin new ice) that were not well represented in the training dataset, were not well captured by the CNN. Hence, there is a sensitivity to the training dataset, which is not surprising given the relatively broad range of conditions represented here, and the limited dataset used.

We suggest that with further, context aware, post processing and training on a more extensive dataset, these results could be made even better. At the same time, they represent a strong step forward in the use of CNNs for automated ice concentration from SAR that does not use manual image analysis charts. While a method has been presented in the literature to estimate SIC from Global Navigation Satellite System Reflectometry data also using a CNN using passive microwave data as training data [38], to our knowledge this is the first doing this for SAR data. This can allow automated use of SAR data in the increasingly important task of monitoring sea ice in high traffic areas of the Canadian Arctic. In addition, the SIC estimated from the CNN is independent of data from atmospheric or ice-ocean models (e.g. ice temperature and wind-speed). Hence, it may be a useful observational dataset for assimilation into coupled ice-ocean or ice-atmosphere data assimilation systems. We suggest that with proper tuning and training, deep convolutional neural networks could become an integral part of large scale ice monitoring.

VIII. ACKNOWLEDGEMENTS

The authors would like to thank Lynn Posgona and Alain Caya at Environment and Climate Change Canada for providing the SAR imagery and image analysis charts. RADARSAT-2 Data and Products MacDonald, copyright Dettwiler and Associates Ltd. 2010. All Rights Reserved. RADARSAT is an official mark of the Canadian Space Agency

REFERENCES

- [1] N. Ivanova, L.T. Pedersen, R.T. Tonboe, S. Kern, T. Lavergne, A. Sorensen, R. Saldo, G. Dybkjaer, L. Brucker, and M. Shokr, “Inter-comparison and evaluation of sea ice algorithms: towards further identification of challenges and optimal approach using passive microwave observations,” *The Cryosphere*, vol. 9, pp. 1797–1817, 2015.
- [2] H. Wiebe, G. Heygster, and T. Markus, “Comparison of the ASI ice concentration algorithm with Landsat-7 ETM+ and SAR imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 9, pp. 3008–3015, 2009.
- [3] K.A. Scott, M. Buehner, and A. Caya, “Direct assimilation of AMSR-E brightness temperatures for estimating

sea ice concentration," *Monthly Weather Review*, vol. 140, pp. 997–1013, 2012.

[4] S. Andersen, R. Tonboe, S. Kern, and H. Schyberg, "Improved retrieval of sea ice concentration from space-borne passive microwave observations using numerical weather prediction model fields: An intercomparison of nine algorithms," *Remote Sensing of Environment*, vol. 104, pp. 374–392, 2006.

[5] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436, 2015.

[6] L. Wang, K.A. Scott, and D.A. Clausi, "Sea ice concentration estimation during melt from dual-pol SAR scenes using deep convolutional neural networks: A case study," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 8, pp. 4524–4533, 2016.

[7] L. Wang, K.A. Scott, and D.A. Clausi, "Sea ice concentration estimation during freeze-up from SAR imagery using a convolutional neural network," *Remote Sensing*, vol. 9, no. 5, pp. doi:10.3390/rs9050408, 2017.

[8] N.G. Kasapoglu, "Synthetic aperture radar data assimilation for sea ice analysis," in *Proceeding of the 2011 International Geoscience and Remote Sensing Symposium*, 2011, pp. 4435–4438.

[9] N.G. Kasapoglu, "Sea Ice Concentration Retrieval Using Composite ScanSAR Features in a SAR Data Assimilation Process," *IEEE Geoscience and Remote Sensing Letters*, vol. 11, no. 12, pp. 2085–2089, 2014.

[10] L. Wang, K. A. Scott, L. Xu, and D. A. Clausi, "Sea ice concentration estimation during melt from dual-pol sar scenes using deep convolutional neural networks: A case study," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 8, pp. 4524–4533, Aug 2016.

[11] L. Wang, K. A. Scott, D. A. Clausi, and Y. Xu, "Ice concentration estimation in the gulf of st. lawrence using fully convolutional neural network," in *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, July 2017, pp. 4991–4994.

[12] J. Kavonen, "Baltic sea ice concentration estimation using SENTINEL-1 and AMSR2 microwave radiometer data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, pp. 2071–2083, 2017.

[13] J. Karvonen, "A sea ice concentration estimation algorithm utilizing radiometer and SAR data," *The Cryosphere*, vol. 8, no. 5, pp. 1639–1650, 2014.

[14] J. Karvonen, "Baltic sea ice concentration estimation based on C-band dual-polarized SAR data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 9, pp. 5558–5566, 2014.

[15] Karvonen J. and M. Mäkinen, "Open water detection from Baltic sea ice Radarsat-1 SAR imagery," vol. 2, no. 3, pp. 275–279, 2005.

[16] H. Han and H. Kim, "Evaluation of summer passive microwave sea ice concentrations in the Chukchi Sea based on KOMPSAT-5 SAR and numerical weather prediction data," *Remote Sensing of Environment*, vol. 209, pp. 343–362, 2018.

[17] A. Komarov and M. Buehner, "Automated detection of ice and open water from dual-polarisation RADARSAT-2 images for data assimilation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, pp. 5755–5769, 2017.

[18] H. Han, S.-H. Hong, H. Kim, T.-B. Chae, and J. Choi, "A study of the feasibility of using KOMPSAT-5 SAR data to map sea ice in the Chukchi Sea in late summer," vol. 8, pp. 467–477, 2017.

[19] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *ICML '09 Proceedings of the 26th Annual International Conference on Machine Learning*, 2009, pp. 41–48.

[20] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez, "Convolutional neural networks for large-scale remote-sensing image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, pp. 645–657, 2017.

[21] A. Kern, S. abd Rosel, L.T. Pedersen, N. Ivanova, R. Saldo, and Tonboe R.T., "The impact of melt ponds on summertime microwave brightness temperatures and sea-ice concentrations," *The Cryosphere*, vol. 10, pp. 2217–2239, 2016.

[22] L. Pogson, T. Geldsetzer, M. Buehner, M. Ross, and K.A. Scott, "Collecting Empirically Derived SAR Characteristic Values over One Year of Sea Ice Environments for Use in Data Assimilation," *Monthly Weather Review*, vol. 145, pp. doi.org/10.1175/MWR-D-16-0110.1, 2017.

[23] Bob Slade, *RADARSAT-2 Product Description*, Nov. 2009.

[24] G. Spreen, L. Kaleschke, and G. Heygster, "Sea ice remote sensing using AMSR-E 89 GHz channels," *Journal of Geophysical Research*, vol. 113, pp. C02S03, 2008.

[25] J. Liu, K.A. Scott, A. Gawish, and P. Fieguth, "Automatic detection of the ice edge in SAR imagery using curvelet transform and active contour," *Remote Sensing*, vol. 8, 2016.

[26] X. Peng, J. Pu, X. Zhao, Q. Ji, and Z. Cheng, "Comparison between AMSR2 sea ice concentration and pseudoship observations of the Arctic and Antarctic sea ice edge on cloud-free days," *Remote Sensing*, vol. 10, pp. doi:10.3390/res10020217, 2018.

[27] G. Crocker, "Canadian ice service digital archive-regional charts: History, accuracy and caveats," *Canadian Ice Service: Ottawa, ON, Canada*, 2006.

[28] *Sea ice analysis and forecasting: Towards an increased reliance on automated prediction systems*, Cambridge University Press, 2017.

[29] G. Huang, Z. Liu, L. v. d. Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 2261–2269.

[30] Yann LeCun, Koray Kavukcuoglu, Clément Farabet, et al., "Convolutional networks and applications in vision," in *ISCAS*, 2010, pp. 253–256.

[31] L. Zhang, Zhang L., and B. Du, "Deep learning for remote sensing data: A technical tutorial on the state of the art," *IEEE Geoscience and Remote Sensing Magazine*, vol. 4, 2016.

- [32] Chris Oliver and Shaun Quegan, *Understanding synthetic aperture radar images*, SciTech Publishing, NC, USOverland, 2004.
- [33] A.S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “CNN features off-the-shelf: an Astounding Baseline for Recognition,” in *Proceeding of the IEEE Computer Science Conference on Computer Vision and Pattern Recognition*, 2014, pp. 512–519.
- [34] X. Yu, Z. Wu, C. Luo, and P. Ren, “Deep learning in remote sensing scene classification: a data augmentation enhanced convolutional neural network framework,” *GI-Science and Remote Sensing*, vol. 54, pp. 741–758, 2017.
- [35] L. Prechelt, “Early Stopping - But When?,” in *Neural Networks: Tricks of the Trade. Lecture Notes in Computer Science*, 7700, pp. 53–67. Springer Berlin Heidelberg, 2012.
- [36] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “ImageNet classification with deep convolutional neural networks.,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [37] K.A. Scott, Z. Ashouri, M. Buehner, T. Carrières, and L. Pogson, “Assimilation of ice and water observations from SAR imagery to improve estimates of sea ice concentration,” *Tellus A*, vol. 67, pp. doi:10.3402/tellusa.v67.27218, 2015.
- [38] Q. Yan and W. Huang, “Sea ice sensing from GNSS data from convolutional neural networks,” *IEEE Geoscience and Remote Sensing Letters*, vol. accepted for publication, pp. doi:10.1109/LGRS.2018.2852143, 2018.



Colin LV Cooke Colin LV Cooke is an undergraduate student pursuing a B.A.Sc. in Mechatronics Engineering (expected April 2019) at the University of Waterloo. He has a strong interest in applied machine learning, explored through various industrial internships and undergraduate researcher assistantships.



K. Andrea Scott K. Andrea Scott received the B.A.Sc. and Ph.D. degrees from the University of Waterloo, Waterloo, ON, Canada, in 1999 and 2008, respectively, and the M.A.Sc. degree from McMaster University, Hamilton, ON, in 2001. She was a Post-Doctoral Researcher with the Data Assimilation and Satellite Meteorology Research Section, Environment and Climate Change Canada, Toronto, ON, where she was part of a team involved in the development of a sea ice data assimilation system. In 2012, she joined the Department of Systems

Design Engineering, University of Waterloo, as a Faculty Member with a specialization in sea ice remote sensing and data assimilation.