

REGION-LEVEL LABELS IN ICE CHARTS CAN PRODUCE PIXEL-LEVEL SEGMENTATION FOR SEA ICE TYPES

Muhammed Patel^{1*} **Xinwei Chen**² **Linlin Xu**¹ **Yuhao Chen**¹
K Andrea Scott³ **David A. Clausi**¹

¹Vision and Image Processing Lab, Systems Design Engineering, University of Waterloo

²School of Marine Science and Engineering, South China University of Technology

³Department of Mechanical and Mechatronics Engineering University of Waterloo

ABSTRACT

Fully supervised deep learning approaches have demonstrated impressive accuracy in sea ice classification, but their dependence on high-resolution labels presents a significant challenge due to the difficulty of obtaining such data. In response, our weakly supervised learning method provides a compelling alternative by utilizing lower-resolution regional labels from expert-annotated ice charts. This approach achieves exceptional pixel-level classification performance by introducing regional loss representations during training to measure the disparity between predicted and ice chart-derived sea ice type distributions. Leveraging the AI4Arctic Sea Ice Challenge Dataset, our method outperforms the fully supervised U-Net benchmark, the top solution of the AutoIce challenge, in both mapping resolution and class-wise accuracy, marking a significant advancement in automated operational sea ice mapping.

1 INTRODUCTION

The increasing impact of rapid climate change on polar sea ice has heightened the importance of monitoring sea ice types using operational Synthetic Aperture Radar (SAR) imagery. This monitoring is crucial for maritime safety, Arctic shipping route planning, and assimilation into forecasting models. SAR satellite observations are commonly used for sea ice mapping, providing wide coverage, day/night observation, high resolution, nearly all-weather capability, and a high revisit rate (Lyu et al., 2022b).

Distinct scattering properties in SAR imagery allow the differentiation of various ice types. Manual annotation of sea ice charts, a common practice for operational sea ice mapping, results in coarse spatial resolution and lacks pixel-level labels, motivating the development of automated methods. Current limitations in manual ice charting, including biases, representativity errors, and labeling uncertainties, underscore the need for more accurate and high-resolution automated sea ice mapping (Stokholm et al. (2022)). In recent years, deep learning (DL) models, such as convolutional neural networks (CNNs), have shown promise due to their ability to automatically extract features, with many related works being published (Boulze et al. (2020); Lyu et al. (2022a); Chen et al. (2023)). However, fully-supervised CNNs require extensive labeled data, which is often limited in sea ice studies.

The AI4Arctic Sea Ice Challenge dataset by Buus-Hinkler et al. (2022) has emerged as a benchmark, providing a large-scale dataset for DL-based sea ice mapping. While this dataset addresses some limitations, challenges remain, such as uncertainties and errors introduced during label conversion.

The ice analyst annotates the SAR image using a polygon-based egg code. This egg code contains the concentration of each ice type within the polygon and lacks pixel-wise labels; therefore, the exact location of the label couldn't be known. However, training a semantic segmentation model requires

*Corresponding author. Email: m32patel@uwaterloo.ca

pixel-wise labels. So, as a workaround, in the AI4Arctic dataset, only dominant ice-type polygons (partial concentration $\geq 65\%$) are used and converted to pixel-wise labels. However, real-world SAR images contain lots of heterogeneous regions, and getting rid of them is not ideal, resulting in the loss of crucial data for training.

To overcome the limitations of fully supervised learning, this paper introduces a novel weakly supervised learning method for per-pixel sea ice type classification directly from region-based ice chart labels, following the successful implementations of weakly supervised learning in remote sensing imagery (Xu et al., 2016). The methodology, experimental results, and conclusions are detailed in subsequent sections.

2 METHODOLOGY

2.1 REGIONAL LABEL GENERATION FROM ICE CHARTS

Despite the SIGRID-3 coding (SIG) in ice charts offering over 10 classes of ice types, some are underrepresented, leading to label uncertainty and visual ambiguities between similar classes. To simplify the problem and enhance generalization, we aggregate SIGRID-3 codes into four classes: open water (ice-free), young ice, first-year ice, and multiyear ice, as described in Table 1. Since the ice characteristics in ice charts are provided for each manually drawn polygon, the regional label for each polygon can be represented by a 4-dimensional vector, indicating the partial fraction of a given ice type. More details about the simplified egg code can be found in Appendix A.2.

Table 1: Description of the classes of sea ice types defined in this research.

Entry number	Class representation	Description
0	Open water	Ice free regions identified in ice charts
1	Young ice	Ice with thickness no more than 30 cm
2	First year ice (FYI)	Ice of no more than one year’s growth developing from young ice with thickness larger than 30 cm
3	Multiyear ice (MYI)	Ice survived at least one melt season




Figure 1: The architecture of the proposed pixel-level ice type mapping model based on a U-Net and the weakly supervised learning scheme (indicated by dash lines). The digit below each layer indicates the number of filters for each convolution/deconvolution layer.

2.2 NETWORK ARCHITECTURE

We employ the U-Net architecture for semantic segmentation (Ronneberger et al., 2015) as the foundation of our ice type classification model, as shown in Fig. 1. The input channels include Sentinel-1 SAR imagery and corresponding multi-source ancillary data, as described in Appendix (Table 4). Other experiment-related hyperparameters, such as input channels, batch size, learning rate, etc., are also outlined in the Appendix A.1.

The U-Net structure in use comprises four encoder-decoder blocks. For the encoder, the number of filters (with 3×3 kernel size) in the 1st to 4th layer is 16, 32, 64, and 64, respectively (and vice versa for the decoder). To predict pixel-based ice type classification, the outputs of the last decoder are processed through a 1×1 convolution layer with the number of filters(channels) equal to the number of classes.

2.3 WEAKLY SUPERVISED LEARNING VIA REGIONAL LOSS REPRESENTATIONS

Our proposed algorithm employs a weakly supervised learning approach, as depicted in Fig. 1. It uses a U-Net architecture to generate a pixel-wise classification map that identifies different types of ice in an input image. During the training phase, we assume each image patch contains an arbitrary number of polygons that have been manually annotated on ice charts. For each of these polygons, we derive a prediction tensor from the U-Net’s output. This prediction tensor, which is a 4-dimensional vector, represents the estimated partial concentrations of various ice classes within the polygon. We then compare this polygon-based prediction tensor against a ground truth tensor derived from the ice chart and calculate the loss. This comparison enables the model to learn from regional (polygon-based) representations, and produce pixel-by-pixel ice types predictions.

Mathematically, if the predicted class probabilities for each pixel in a polygon are denoted as $\mathbf{y} = \{y_1, y_2, y_3, y_4\}$, then each element in the 4 dimensional polygon-based predicted tensor (denoted as \mathbf{Y}) can be calculated as

$$Y_i = \frac{1}{N_{pixel}} \sum_{j=1}^{N_{pixel}} y_i^j, \quad i = 1, \dots, 4, \quad (1)$$

where N_{pixel} refers to the total number of pixels inside a polygon, and y_i^j represents y_i for the j^{th} pixel. In this work, we use the cross-entropy (CE) loss, expressed as

$$\text{loss} = -\frac{1}{n} \sum_{i=1}^n Y_i \log(\hat{Y}_i), \quad (2)$$

where n is the number of elements (classes) in each predicted/ground truth polygon-based tensor, and Y_i and \hat{Y}_i are the ground truth and predicted concentrations for the i^{th} class from the polygon, respectively. Finally, the loss calculated from all the polygons in all image patches in a batch is aggregated as the total loss for backpropagation. By minimizing this total loss using optimization algorithms (e.g., stochastic gradient descent Loshchilov & Hutter (2016)), the model effectively learns the distribution of different ice types in each polygon while accurately providing per-pixel classification. The pseudocode for this loss function can be found in Appendix A.3

3 EXPERIMENTAL RESULTS

We perform our experiments on the AI4arctic dataset and follow the same train-test split as in the Autoice competition ((Stokholm et al., 2023)) for direct comparison. Out of the 533 dataset files, 22 are reserved for testing, and cross-validation is performed on the training set. Each run involves randomly selecting 20 scenes for validation, with a total of 30 runs ensuring extensive scene coverage.

Given the absence of pixel-based ground truth, we numerically evaluate the model by computing R2 scores between polygon-based sea ice type labels and predictions representing partial concentrations. The R2 score for a specific class R_i^2 is calculated using the formula:

$$R_i^2 = 1 - \frac{\sum_{k=1}^{N_{poly}} (Y_i^k - \hat{Y}_i^k)^2}{\sum_{k=1}^{N_{poly}} (Y_i^k - \bar{Y}_i)^2}, \quad (3)$$

where Y_i^k and \hat{Y}_i^k are the ground truth and predicted partial concentration for the i^{th} class in the k^{th} polygon, \bar{Y}_i is the mean value of Y_i^k across all polygons, and N_{poly} is the total number of polygons in the validation/testing set.

We compare our weakly supervised U-Net with the winner of the AutoIce competition, referred to as benchmark U-Net (Chen et al. (2024)). We keep all the hyperparameters (A.1) same and only

Table 2: The class-wise R2 scores calculated between predictions and ice charts of each ice type using the benchmark U-Net and the proposed weakly supervised U-Net. The improvement of R2 scores for each class using the weakly supervised model are indicated by the percentages in the brackets.

Models	Cross validation				Testing			
	Open water	Young ice	First year ice	Multiyear ice	Open water	Young ice	First year ice	Multiyear ice
Benchmark U-Net trained with per pixel labels	93.95%	57.57%	75.54%	80.78%	91.37%	41.22%	75.52%	83.61%
Weakly supervised U-Net	97.45% (+4.00%)	68.16% (+10.59%)	85.23% (+9.69%)	88.93% (+8.17%)	95.73% (+4.36%)	58.83% (+17.61%)	83.09% (+7.57%)	86.04% (+2.43%)

replace the loss function and ground truth used. The benchmark model is trained with pixel-wise labels and cross-entropy loss. Table 2 displays class-wise R2 scores computed between predictions and ice charts labels on validation and testing sets. The weakly supervised U-Net achieves higher R2 scores in all ice types. In the cross-validation sets, it increases R2 scores by around 4%, 10.59%, 9.69%, and 8.17% for open water, young ice, first-year ice, and multiyear ice, respectively. The same improvements are observed in the testing set: 4.36%, 17.61%, 7.57%, and 2.43%. The hypothesis for these improvements is that the benchmark U-Net only uses pixel-wise labels derived from dominant polygons and ignores mixed polygons. This gets rid of so many polygons, which usually contain minority classes (Young ice, First-year ice). On the other hand, our method is more aligned with the way the labels are defined in charts.

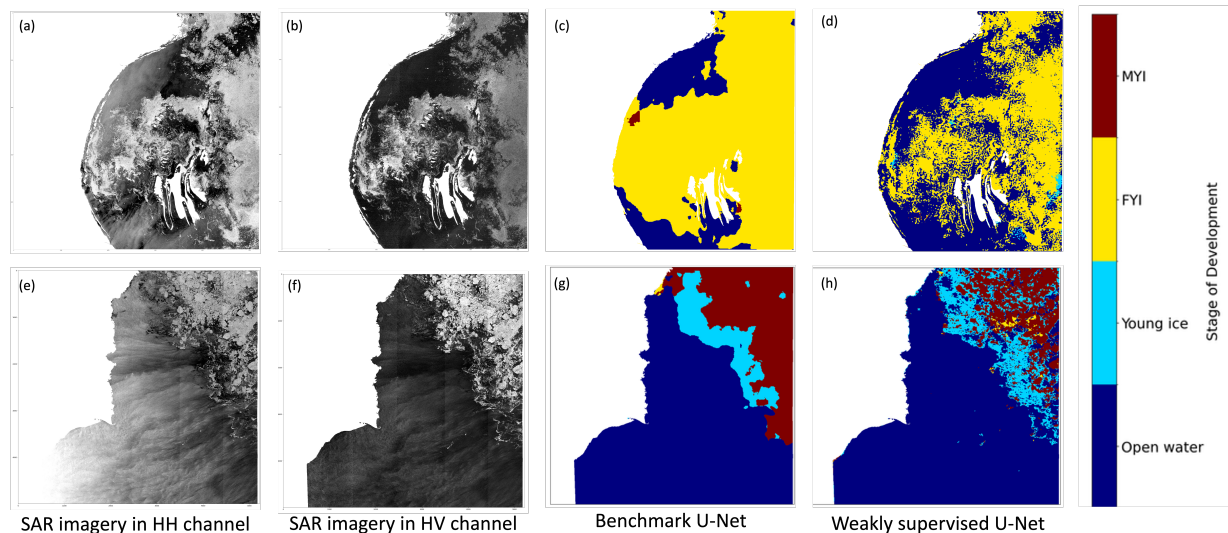


Figure 2: Visualization of results obtained from two SAR scenes. (a): SAR image in HH channel, file ID name: 20180707T113313_cis, collected on July 7th, 2018 in Hudson Bay (center latitude/longitude: 57.07°N, 77.96°W); (b): The HV channel of the SAR image in (a); (c): Classification results of (a) from the benchmark U-Net model (dark blue: open water; light blue: young ice; yellow: first year ice; red: multiyear ice); (d): Classification results of (a) from the proposed weakly supervised U-Net; (e): SAR image in HH channel, file ID name: 20190926T152005_cis, collected on September 26th, 2019 in Western Arctic (center latitude/longitude: 72.63°N, 126.61°W); (f): The HV channel of the SAR image in (e); (g): Classification results of (e) from the benchmark U-Net model; (h): Classification results of (e) from the proposed weakly supervised U-Net. The land area is masked in white.

While numerical results offer insights into accuracy, visual analysis in Fig. 2 provides a more qualitative assessment. The weakly supervised model significantly enhances mapping resolution, particularly in identifying ice-water boundaries, polygon-based sea ice concentration, and partial concentrations of different ice types. This improvement is evident in challenging seasons, such as

melt and freeze-up periods. For example, in the SAR scene on the top of Fig. 2 (melt period), the weakly supervised model is able to detect the detailed ice-water boundaries and ice with small floe size surrounded by water precisely, while the benchmark model can only detect ice or water-dominated areas with coarse resolution. For the SAR scene at the bottom obtained in freeze-up, pixels dominated by young ice (with relatively low backscatter intensities) and multiyear ice (with very bright pixel intensities) are well separated with significantly enhanced resolution using the weakly supervised model. Additional visualization with zoom-in example can be found in Appendix A.4

4 CONCLUSION

We introduce an automated pipeline for sea ice type classification using SAR imagery and ancillary data, leveraging a weakly supervised U-Net. This approach holds potential for integration into operational sea ice mapping for tactical navigation. By enabling precise prediction of sea ice type distributions with high resolution, our method eliminates the need for per-pixel labels during model training. To address the challenge of limited high-resolution ground truths, our classification model employs weakly supervised learning techniques, incorporating regional labels and loss representations. This allows the utilization of region labels derived from ice charts with coarse resolution directly for model training. The model is trained and validated using the recently released AI4Arctic Sea Ice Challenge Dataset. Results demonstrate strong agreement between predictions and ice-type distributions in ice charts, validated through visual interpretation. In future works, we'll aim to test our method on other datasets and state-of-the-art segmentation models.

5 ACKNOWLEDGEMENT

The authors acknowledge the providers of the AI4Arctic dataset and the organizers of the AutoIce competition. This work was supported in part by Environment and Climate Change Canada (ECCC) and the Natural Sciences and Engineering Research Council (NSERC) of Canada. This research was enabled in part by support provided by Calcul Québec's (calculquebec.ca) and the Digital Research Alliance of Canada (alliancecan.ca).

REFERENCES

- IICWG: SIGRID-3: a vector archive format for sea ice charts. <https://library.wmo.int/records/item/37171-sigrid-3-a-vector-archive-format-for-sea-ice-charts>. Accessed: 2023-10-20.
- Hugo Boulze, Anton Korosov, and Julien Brajard. Classification of sea ice types in Sentinel-1 SAR data using convolutional neural networks. *Remote Sens.*, 12(13):2165, 2020. doi: 10.3390/rs12132165.
- Jørgen Buus-Hinkler, Tore Wulf, Andreas Rønne Stokholm, Anton Korosov, Roberto Saldo, Leif Toudal Pedersen, and et al. AI4Arctic Sea Ice Challenge Dataset. Technical University of Denmark. Collection, 2022. URL <https://doi.org/10.11583/DTU.c.6244065.v2>.
- Xinwei Chen, K Andrea Scott, Mingzhe Jiang, Yuan Fang, Linlin Xu, and David A Clausi. Sea ice classification with dual-polarized SAR imagery: A hierarchical pipeline. In *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, pp. 224–232, 2023.
- Xinwei Chen, Muhammed Patel, Fernando J Pena Cantu, Jinman Park, Javier Noa Turnes, Linlin Xu, K Andrea Scott, and David A Clausi. Mmseaiice: a collection of techniques for improving sea ice mapping with a multi-task model. *The Cryosphere*, 18(4):1621–1632, 2024.
- Joint WMO-IOC Technical Commission for Oceanography and Marine Meteorology. Ice chart colour code standard, version 1.0, 2014. Technical Report (WMO TD: 1215b), World Meteorological Organization & Intergovernmental Oceanographic Commission, Geneva, Switzerland, 2014.
- Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. URL <https://arxiv.org/abs/1608.03983>.
- Hangyu Lyu, Weimin Huang, and Masoud Mahdianpari. Eastern arctic sea ice sensing: First results from the RADARSAT Constellation Mission data. *Remote Sens.*, 14(5):1165, 2022a.
- Hangyu Lyu, Weimin Huang, and Masoud Mahdianpari. A meta-analysis of sea ice monitoring using spaceborne polarimetric SAR: Advances in the last decade. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, 15:6158–6179, 2022b.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Intl. Conf. Med. Image Comput. Comput. Assist. Interv.*, pp. 234–241. Springer, 2015. doi: 10.1007/978-3-319-24574-4_28.
- Andreas Stokholm, Tore Wulf, Andrzej Kucik, Roberto Saldo, Jørgen Buus-Hinkler, and Sine Munk Hvidegaard. AI4Sealce: Toward solving ambiguous SAR textures in convolutional neural networks for automatic sea ice concentration charting. *IEEE Trans. Geosci. Remote Sens.*, 60:1–13, 2022.
- Andreas Stokholm, Jørgen Buus-Hinkler, Tore Wulf, Anton Korosov, Roberto Saldo, David Arthurs, Rune Solberg, Nicolas Longépé, and Matilde Kreiner. The AutoICE competition: Automatically mapping sea ice in the Arctic. Technical report, Copernicus Meetings, 2023.
- Linlin Xu, David A Clausi, Fan Li, and Alexander Wong. Weakly supervised classification of remotely sensed imagery using label constraint and edge penalty. *IEEE Trans. Geosci. Remote Sens.*, 55(3):1424–1436, 2016.

A APPENDIX

A.1 MODEL HYPERPARAMETERS

Table 3: Specifications of training the proposed model.

Optimizer	Stochastic gradient descent with momentum (SGDM)
Learning rate	0.001
Weight decay	0.01
Scheduler	Cosine Annealing
Batch size	16
Number of iterations per epoch	500
Total epoch	50
Number of epochs for the first restart	50
Image Downscaling ratio	10
Patch size	256

Table 4: The input channels of the proposed model.

Input type	Number of channels	Description
SAR imagery	2	Dual polarized Sentinel-1 SAR images in HH and HV channels
Passive microwave	2	AMSR2 brightness temperature data in 36.5 GHz in both horizontal and vertical polarizations
Spatial information	2	The latitude and longitude of each pixel in the SAR imagery
Temporal information	1	The acquisition month of the SAR imagery

A.2 SIMPLIFIED EGGCODE

An overview of the ice egg code variables in the simplified egg code is presented in Fig. 3 (a). Compared to the original egg code in WMO standard (Joint WMO-IOC Technical Commission for Oceanography and Marine Meteorology, 2014), the information concerning the form of ice (floe size) is omitted. The total concentration, C_t refers to the total SIC (Sea Ice concentration) of the polygon. As C_a , C_b , and C_c correspond to the partial concentrations of the top 3 thickest ice types, there is a relation that $C_a + C_b + C_c = C_t$. On the other hand, S_a , S_b , and S_c are represented by the entry numbers in Table 1, so for e.g. in the most left egg code in Fig. 3(b), the value 9 represents C_t , values 2, 7, 0 denotes partial concentration C_a , C_b , C_c respectively. Similarly, 3, 2, 0 represents the class S_a , S_b , S_c . And the final ground-truth 4 dimensional vector would be $Y_i = \{1, 2, 7, 0\}$ where the 1st element represents the concentration of water which is nothing but $1 - C_t$.

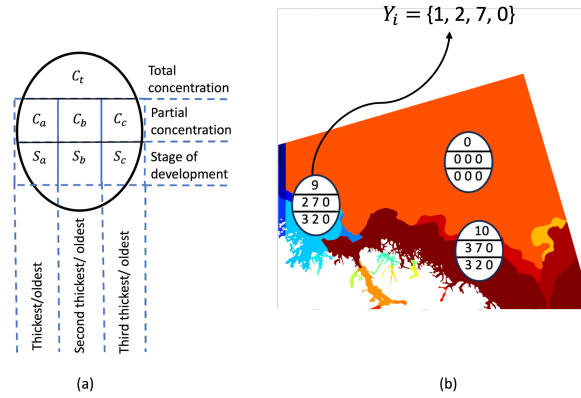


Figure 3: (a) Overview of the ice egg code parameters in the simplified egg code produced from the original ice charts. Concentration values range from 0 (empty) to 10 (fully- covered) with each increase of 1 representing a 10 percent concentration increase. (b) An example Ice chart showing eggcode and conversion of eggcode to ground truth tensor.

A.3 LOSS FUNCTION: PSUEDOCODE

Algorithm 1: Polygon-based Loss Calculation Pseudocode based on Python

Data: polygon_labels, pixel_predictions, land_masks, polygon_maps

Result: total_loss

total_loss = 0

batch_size = pixel_predictions.shape[0]

for *img_id* *in* range(*batch_size*) **do**

 polygon_ids = np.unique(polygon_maps[*img_id*]);

for *polygon_id* *in* *polygon_ids* **do**

 // get the label tensor for a certain polygon (*polygon_id*)

polygon_label = *polygon_labels*[*img_id*][*polygon_id*]

 // exclude the polygons that are masked out from loss calculation

if *polygon_label* == 255 **then**

 loss = 0

 total_loss += loss

continue;

 // create a mask to mask out pixels outside the current polygon

poly_mask = (*polygon_icecharts*[*img_id*] == *poly_id*)

 // create final mask by removing land area from the *poly_mask*

final_mask = *poly_mask* ∩ *land_masks*[*img_id*]

 // pixel-based SOD predictions within the current polygon

pixel_AOI = *pixel_predictions*[*img_id*][*final_mask*]

 // aggregate and average over pixel-based predictions to generate polygon-based prediction

polygon_prediction = *pixel_AOI*.sum(0)/*pixel_AOI*[0].numel()

 // calculate Cross Entropy loss between label tensor and predicted tensor

 loss = Cross_Entropy(*polygon_prediction*, *polygon_label*)

 total_loss += loss

return *total_loss*

A.4 ADDITIONAL VISUALIZATION

The zoom-in examples from Fig. 2 are provided in Fig 4. In the first column of Fig 4, the weakly supervised model is able to detect the detailed ice-water boundaries and ice with small floe size surrounded by water precisely, while the benchmark model can only detect ice or water-dominated areas with coarse resolution. In the second column of Fig. 4, pixels dominated by young ice (with relatively low backscatter intensities) and multiyear ice (with very bright pixel intensities) are well separated with each other using the weakly supervised model.

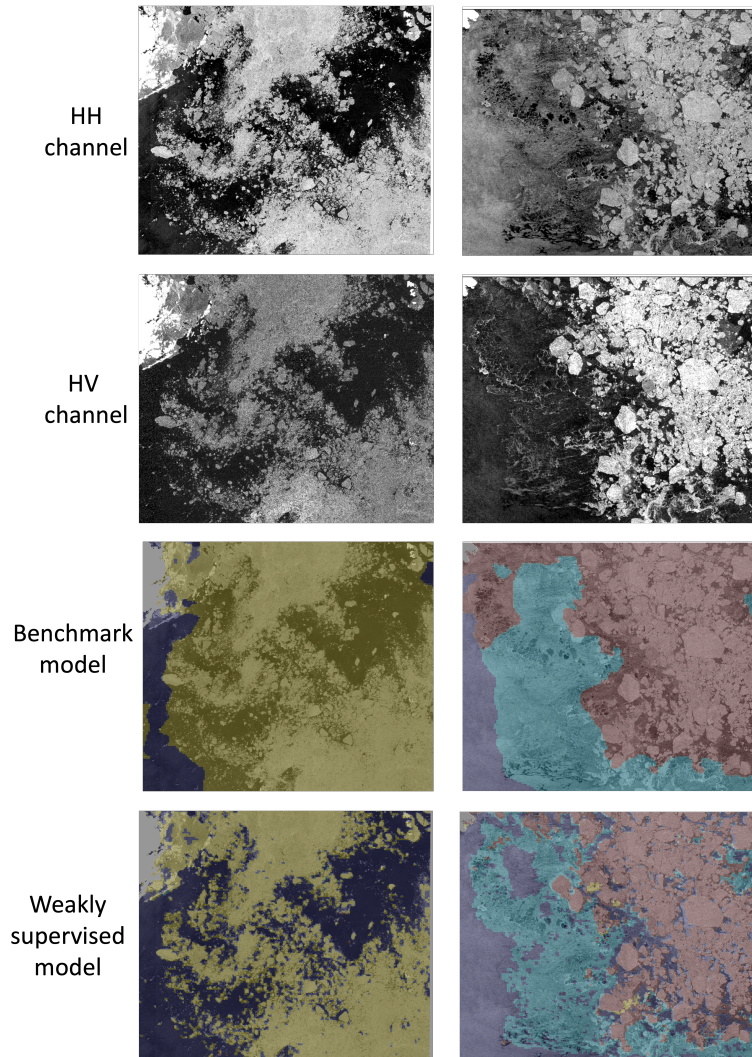


Figure 4: Zoom-in examples of regions cropped from SAR scenes and corresponding predictions in Fig. 2. First column: a region located within a polygon in Fig. 2 dominated by first-year ice. Second column: a region located within a polygon in Fig. 2 consisting of both young ice and multiyear ice. First row: regions in HH channel; Second row: regions in HV channel; Third row: overlay of HH channel and predictions from benchmark model, with the color bar provided in Fig. 2; Fourth row: overlay of HH channel and predictions from weakly supervised model.