

Real-time GeoAI for High-resolution Mapping and Segmentation of Arctic Permafrost Features

The case of ice-wedge polygons

Wenwen Li*

School of Geographical Sciences and Urban Planning, Arizona State University
Tempe, AZ, USA
wenwen@asu.edu

Chia-Yu Hsu

School of Geographical Sciences and Urban Planning, Arizona State University
Tempe, AZ, USA
chsu53@asu.edu

Sizhe Wang

School of Computing and Augmented Intelligence, Arizona State University
Tempe, AZ, USA
wsizhe@asu.edu

Chandi Witharana

Department of Natural Resources and the Environment, University of Connecticut
Storrs, CT, USA
chandi.witharana@uconn.edu

Anna Liljedahl

Woodwell Climate Research Center
Falmouth, MA, USA
aliljedahl@woodwellclimate.org

ABSTRACT

This paper introduces a real-time GeoAI workflow for large-scale image analysis and the segmentation of Arctic permafrost features at a fine-granularity. Very high-resolution (0.5m) commercial imagery is used in this analysis. To achieve real-time prediction, our workflow employs a lightweight, deep learning-based instance segmentation model, SparseInst, which introduces and uses Instance Activation Maps to accurately locate the position of objects within the image scene. Experimental results show that the model can achieve better accuracy of prediction at a much faster inference speed than the popular Mask-RCNN model.

CCS CONCEPTS

- Computing methodologies → Neural networks; Image segmentation.

KEYWORDS

GeoAI, Artificial Intelligence, Arctic, Permafrost, Instance segmentation

ACM Reference Format:

Wenwen Li, Chia-Yu Hsu, Sizhe Wang, Chandi Witharana, and Anna Liljedahl. 2022. Real-time GeoAI for High-resolution Mapping and Segmentation of Arctic Permafrost Features: The case of ice-wedge polygons. In *Proceedings of ACM SIGSPATIAL International Conference on Advances in*

*Corresponding Author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGSPATIAL '22, November 1-4, 2022, Seattle, WA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9529-8/22/11...\$15.00

<https://doi.org/10.1145/3557918.3565869>

Geographic Information Systems (SIGSPATIAL '22). ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3557918.3565869>

1 INTRODUCTION

Polar regions are one of Earth's remaining frontiers that play a vital role in global climate, ecosystems, and economy. Global warming over the past century is driving dramatic change in the Arctic ecosystem, endangering its natural environment, infrastructure, and life of the indigenous population. Permafrost, ground that remains below 0°C for at least two consecutive summers, is at the center of this change. Covering nearly $\frac{1}{4}$ of the land in the northern hemisphere, thawing permafrost is causing significant local and regional impacts on the Arctic community. As the ice-rich frozen ground thaws, land subsides causing severe damage to buildings, roads, pipelines, and industrial infrastructure [9]. Permafrost degradation also increases rates of coastal erosion, wildfires, and flooding, which may further accelerate the thawing process and make the Arctic ecosystem even more vulnerable to climate change [6]. At a global scale, the thawing of Arctic permafrost will result in the release of an immense amount of carbon dioxide and methane, exaggerating the greenhouse effect and global warming through complex feedback mechanisms [17].

To improve our understanding of permafrost dynamics and its linkages to other Arctic ecosystem components in the midst of rapid Arctic change, it is critically important to have spatial data readily available that provide fine-granularity mapping of permafrost features, their extent, distribution, and longitudinal changes. Achieving this goal requires new approaches that can perform automated mining from Arctic big data. It is exciting that the Arctic community has started to embrace GeoAI [11, 12] and big data to support Arctic research, from predicting Arctic sea ice concentration [1], to finding marine mammals on ice [15], creating Arctic land cover maps [18], and automated mapping of permafrost features [2]. Pioneering research in performing automated characterization of Arctic permafrost features has also been reported in the literature. An

GeoAI-based Mapping Application for Permafrost Land Environment (MAPLE) is being developed to integrate Big Imagery, GeoAI, and High-Performance Computing (HPC) to achieve classification of permafrost features, in particular, ice-wedge polygons (IWP) [19]. The delineation of IWPs is achieved using a popular instance segmentation model, Mask R-CNN [7]. Huang et al. [10] applied a semantic segmentation model U-Net for mapping retrogressive thaw slumps, another important feature type of Arctic permafrost for understanding permafrost thaw and Arctic warming.

While these deep learning models, such as Mask R-CNN, result in satisfying performance in terms of prediction accuracy, they can hardly achieve real-time processing because the algorithms often require placement of a large number of candidate bounding boxes and complex post-processing to remove redundant information. To reduce computational cost and perform efficient permafrost mapping at the pan-Arctic scale (which covers over 5 million km² of tundra region), it is necessary to develop and apply new models that can achieve high-accuracy and real-time prediction. This paper aims to achieve this goal by integrating a novel real-time instance segmentation model, SparseInst [5], in our automated permafrost feature mapping pipeline. The next section describes the methodological workflow in detail.

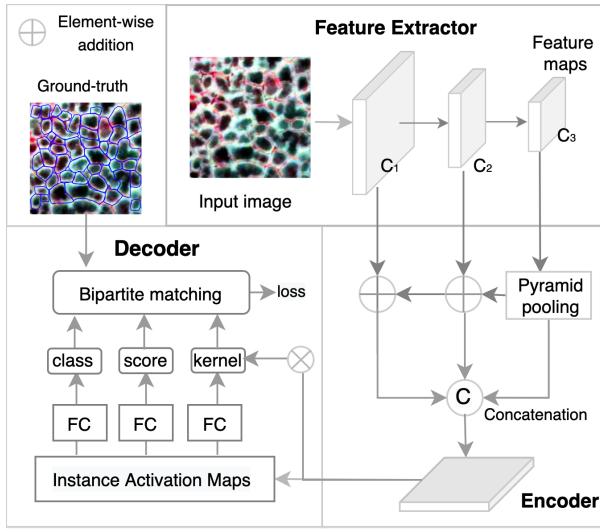


Figure 1: Real-time GeoAI workflow for Arctic permafrost segmentation and mapping. FC: Fully Connected layer

2 METHOD

Figure 1 demonstrates the workflow of real-time GeoAI for Arctic permafrost mapping. We adopt a novel instance segmentation model SparseInst into the workflow, which contains three major components: a feature extractor, an instance context encoder, and an Instance Activation Map (IAM)-based decoder. The feature extractor is responsible for extracting multi-scale features from the input. The encoder will process the extracted features and fuse them into single-level features with multi-scale representations. The encoded features are then processed by the decoder to generate IAMs for instance classification and segmentation. Each component

is designed under the consideration of lightweight architecture and low computational complexity to achieve fast inference speed.

2.1 Feature Extractor

The feature extractor adopted in this work is ResNet-50 [8]. Among various deep neural network (DNN) architectures, ResNet-50 enjoys a good trade-off between accuracy and model complexity so to support real-time applications [3]. ResNet extracts representative features for objects of different types using a deep residual network. After a series of convolutional operations, multi-scale feature maps can be generated, among which high-resolution maps are better at small-object segmentation and low-resolution feature maps can better support segmentation of large objects. To accurately segment objects of varying sizes, hierarchical feature maps at multiple scales and resolutions are passed to the encoder (see Figure 1).

2.2 Instance Context Encoder

The main purpose of the encoder is to generate a single feature map containing multi-scale representations. Conventional approaches use multi-scale features with multi-level predictions [13] for segmenting objects at different scales [20]. However, this will increase overall processing time of the model, making it less efficient and less favorable for real-time applications. Recent real-time instance segmentation models [4, 16] fuse multi-scale information into a single feature map to reduce both prediction and post-processing time. SparseInst utilizes a similar idea and it fuses three feature maps obtained from different convolution stages. The fusion first follows the feature pyramid network (FPN) [14] to use a top-down pathway for building semantic-rich features. To further enhance the scale information, the last feature map (C₃) also undergoes a pyramid pooling operation [23] to increase the global contextual information without increasing the size of the feature maps. Next, all feature maps are upsampled to the same resolution and concatenated together to generate feature maps at a single resolution but with multi-scale representations. The output is then sent to the decoder for classification and segmentation.

2.3 IAM-based Decoder

The function of the decoder is to take the fused feature map from the encoder as input to generate N predictions. Each prediction contains a triple <object class, objectness score, kernel>. The objectness score refers to the probability of an object belonging to a certain class and the kernel is a low-dimensional representation of location information for that object. This instance-level prediction is achieved through the generation of Instance Activation Maps (IAMs) which are capable of highlighting important image areas. Different from conventional approaches which use dense anchors to detect and segment objects, SparseInst trains the decoder to create IAMs, which have a one-to-one mapping with the objects to segment. This design helps the decoder to achieve real-time performance as it avoids the time-consuming post-processing of some models, such as Mask R-CNN, which need to select from thousands of anchors to predict the most accurate mask and to perform matching between predicted masks and the ground-truth. Once the predictions are generated, they are sent to perform bipartite matching to associate each ground-truth object with its most

Table 1: Comparisons with Mask R-CNN [7] for mask AP and speed on IWP dataset. Inference speeds of all models are tested with single NVIDIA A5000 GPU.

Model	FPS	AP ₅₀	AP _S	AP _M	AP _L
Mask R-CNN	27.01	52.86	33.28	60.03	64.39
SparseInst	45.61	53.97	31.70	60.78	68.10

similar prediction, then the difference between the prediction and the ground-truth is encoded into the loss function. As the model is being trained, it learns to generate more accurate IAMs and thus more accurate predictions, lowering the loss until the model fully converges.

3 EXPERIMENTS AND RESULTS

3.1 Data

To assess the performance of the models, we created an AI-ready dataset containing 867 image tiles and a total of 34,931 ice-wedge polygons (IWPs). The dataset covers dominant tundra vegetation types in the polygonal landscapes, including sedge, tussock, and barren tundra. Very high resolution (0.5 m) remote sensing imagery acquired by Maxar sensors is used for annotation and model training. The average image size is $\sim 226 \times 226$ with the largest image size 507×507 . Each image has a label indicating the image size and coordinates of the IWPs.

The labeled images are divided into three sets: training (70%), validation (15%), and testing (15%). The maximum number of IWPs per image is 447. This statistic is critical in determining the maximum number of detections per image, as it is an important hyperparameter to set in the segmentation model. It also affects both accuracy and speed and provides a trade-off between them (Section 3.3).

3.2 Model Training and Results

In this work, we compare SparseInst with one of the most popular instance segmentation models, Mask R-CNN [7]. Both models are built upon Detectron2 [21], a module of the PyTorch deep learning framework which provides state-of-the-art segmentation algorithms. The training is conducted on four NVIDIA A5000 GPUs. The batch size is 16 and the maximum number of iterations is 20,000. The maximum number of detections per image N is set to 500. Table 1 shows the performance comparison between Mask R-CNN (default setting) and SparseInst. The evaluation metric for model inference speed is frame per second (FPS) and for accuracy, average precision (AP) [22] is used. As the results show, SparseInst demonstrates better performance in terms of both speed and accuracy than Mask R-CNN. We also separate IWPs into three groups by their areas: small (area < 200 pixels), medium (area in between 200 and 450 pixels), and large (area > 450 pixels). Table 1 also shows the average precision (AP) in each group. SparseInst performs slightly worse than Mask R-CNN on small IWPs segmentation, but it works better at segmenting medium- to large-size IWPs. Overall, SparseInst yields better detection accuracy than Mask R-CNN. Speed-wise, the model runs nearly twice as fast as Mask R-CNN, achieving real-time performance (model's inference speed at 30 FPS or above).

3.3 Precision vs. Speed

Figure 2 shows the precision and speed trade-off of the SparseInst model and its comparison with Mask R-CNN. We used the default setting of Mask R-CNN to conduct training and testing as it achieves better performance than other experimental settings. Differently, SparseInst requires a predefined N to determine the maximum number of masks and predictions per image. This hyperparameter not only affects the model's prediction accuracy but also its speed. A larger N will slow down the process of bipartite matching during training and increase model complexity in the decoder part, therefore negatively affecting the model's efficiency during both training and testing. Here, we tested the model performance at different settings of N (at 100, 300 and 500 respectively). It can be seen that as N decreases, the model's prediction speed increases (x axis) but its predictive power (y axis) decreases (from 54% at $N=500$ to 51% at $N=100$). For Mask R-CNN, while its prediction accuracy is quite high, the speed is below the threshold of models that can be considered real-time. It is noteworthy that at both $N=500$ and $N=300$, SparseInst achieves better prediction accuracy than Mask R-CNN. This result verifies the importance of carefully setting values of hyperparameters according to data characteristics to achieve satisfying model performance.

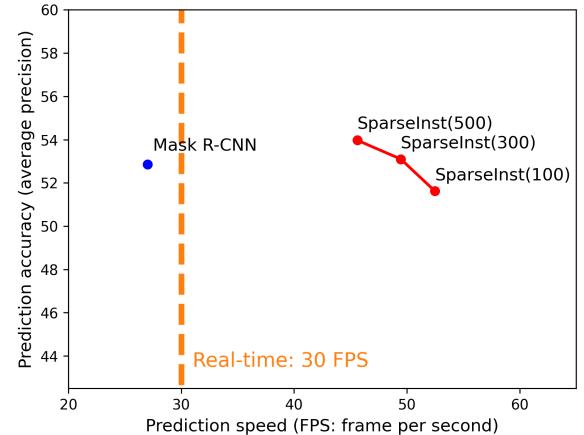


Figure 2: Speed and accuracy trade-off.

3.4 Prediction Results

Figure 3 illustrates segmentation results for two sample images. Figure 3a and 3c provides the ground-truth labels of the IWPs. The ice-wedge polygons in these two images belong to two distinctive types of IWPs: low-centered (3a, 3b) and high-centered (3c, 3d). A preliminary analysis has also shown that when separating these feature types, thus making the segmentation task more challenging, the performance advantage of SparseInst over Mask R-CNN become even more dominant. This reflects the robustness of the SparseInst model in performing high-accuracy and real-time IWP segmentation.

Figure 3b and 3d present the model prediction results for the two images to their left (3a and 3c). It can be seen that for smaller objects, although the predicted area is quite close to the ground-truth, the

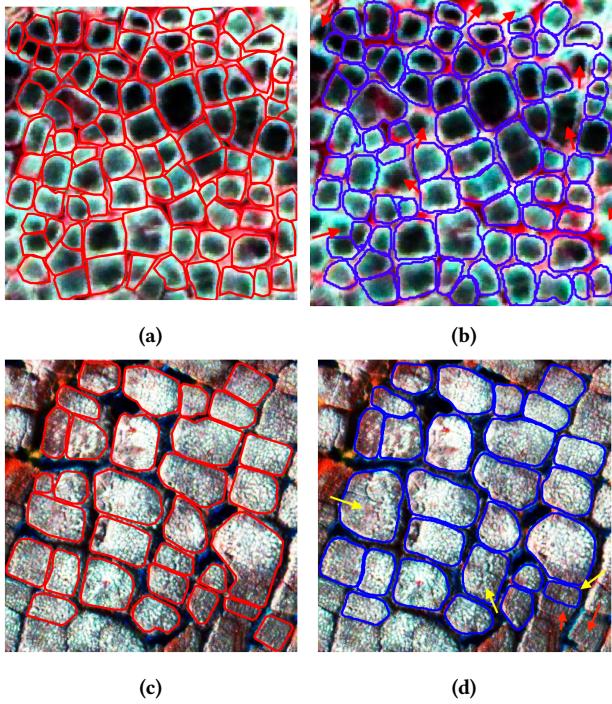


Figure 3: Comparison between ground-truth (a and c) and model segmentation results (b and d). Red arrow: missing prediction; Yellow arrow: incorrect prediction

boundary line itself is not as smooth as the human labels, 3b. This issue does not exist in segmentation results for large objects, 3d. The model did miss predictions for a few IWP when there exist no clear boundaries around them (red arrows in 3b and 3d). There are also incorrect predictions (yellow arrows in 3d); this is likely due to the semantically different concepts that expert annotators and the machine consider. Interestingly, the model can predict labels for some partial IWP near the border where is not labeled by experts.

4 CONCLUSION

This paper introduces a real-time GeoAI workflow for segmenting an important permafrost feature, IWP. Delineating their extent and qualifying their changes is critically important to understand Arctic warming and permafrost thaw and its impact to the Arctic environment, infrastructure, and people. Here, we adopt a light-weight instance segmentation model into the workflow and verify its good performance in terms of both prediction accuracy and speed. In the future, we will further improve both the training data to explicitly annotate multi-type IWP, and also refine the model to improve its detection accuracy of small objects.

ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation under awards 2120943, 1853864, 1927872, and 2230034.

REFERENCES

- [1] Tom Andersson, Fruzsina Agocs, Scott Hosking, María Pérez-Ortiz, Brooks Paige, Chris Russell, Andrew Elliott, Stephen Law, Jeremy Wilkinson, Yevgeny Askenov, et al. 2020. Deep learning for monthly Arctic sea ice concentration prediction. In *EGU General Assembly Conference Abstracts*. 15481.
- [2] Md Abul Ehsan Bhuiyan, Chandi Witharana, and Anna K Liljedahl. 2020. Use of very high spatial resolution commercial satellite imagery and deep learning to automatically map ice-wedge polygons across tundra vegetation types. *Journal of Imaging* 6, 12 (2020), 137.
- [3] Simone Bianco, Remi Cadene, Luigi Celona, and Paolo Napoletano. 2018. Benchmark Analysis of Representative Deep Neural Network Architectures. *IEEE access* 6 (2018), 64270–64277.
- [4] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. 2019. Yolact: Real-time Instance Segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9157–9166.
- [5] Tianheng Cheng, Xinggang Wang, Shaoyu Chen, Wenqiang Zhang, Qian Zhang, Chang Huang, Zhaoxiang Zhang, and Wenyu Liu. 2022. Sparse Instance Activation for Real-Time Instance Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4433–4442.
- [6] Guido Grosse, Scott Goetz, A Dave McGuire, Vladimir E Romanovsky, and Edward AG Schuur. 2016. Changing permafrost in a warming world and feedbacks to the Earth system. *Environmental Research Letters* 11, 4 (2016), 040201.
- [7] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [9] Jan Hjort, Olli Karjalainen, Juha Aalto, Sebastian Westermann, Vladimir E Romanovsky, Frederick E Nelson, Bernd Etzelmüller, and Miska Luoto. 2018. Degrading permafrost puts Arctic infrastructure at risk by mid-century. *Nature communications* 9, 1 (2018), 1–9.
- [10] Lingcao Huang, Trevor C Lantz, Robert H Fraser, Kristy F Tiampo, Michael J Willis, and Kevin Schaefer. 2022. Accuracy, Efficiency, and Transferability of a Deep Learning Model for Mapping Retrogressive Thaw Slumps across the Canadian Arctic. *Remote Sensing* 14, 12 (2022), 2747.
- [11] Wenwen Li. 2020. GeoAI: Where machine learning and big data converge in GIScience. *Journal of Spatial Information Science* 20 (2020), 71–77.
- [12] Wenwen Li and Chia-Yu Hsu. 2022. GeoAI for Large-Scale Image Analysis and Machine Vision: Recent Progress of Artificial Intelligence in Geography. *ISPRS International Journal of Geo-Information* 11, 7 (2022), 385.
- [13] Yanghao Li, Yuntao Chen, Naiyan Wang, and Zhao-Xiang Zhang. 2019. Scale-Aware Trident Networks for Object Detection. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Seoul, Korea (South), 6053–6062. <https://doi.org/10.1109/ICCV.2019.00615>
- [14] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature Pyramid Networks for Object Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2117–2125.
- [15] NOAA. 2019. Developing artificial intelligence to find ice seals and polar bears from the sky. <https://www.fisheries.noaa.gov/feature-story/developing-artificial-intelligence-find-ice-seals-and-polar-bears-sky> (2019).
- [16] Hughes Perreault, Guillaume-Alexandre Bilodeau, Nicolas Saunier, and Maguelonne Héritier. 2021. CenterPoly: Real-Time Instance Segmentation Using Bounding Polygons. *arXiv:2108.08923 [cs]* (Sept. 2021). [arXiv:2108.08923 \[cs\]](https://arxiv.org/abs/2108.08923)
- [17] Edward AG Schuur, A David McGuire, Christina Schädel, Guido Grosse, Jennifer W Harden, Daniel J Hayes, Gustaf Hugelius, Charles D Koven, Peter Kuhry, David M Lawrence, et al. 2015. Climate change and the permafrost carbon feedback. *Nature* 520, 7546 (2015), 171–179.
- [18] Hunsoo Song, Yonghyun Kim, and Yongil Kim. 2019. A patch-based light convolutional neural network for land-cover mapping using Landsat-8 images. *Remote Sensing* 11, 2 (2019), 114.
- [19] Mahendra R Udalwalpola, Amit Hasan, Anna Liljedahl, Aiman Soliman, Jeffrey Terstriep, and Chandi Witharana. 2022. An Optimal GeoAI Workflow for Pan-Arctic Permafrost Feature Detection from High-Resolution Satellite Imagery. *Photogrammetric Engineering & Remote Sensing* 88, 3 (2022), 181–188.
- [20] Xilong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. 2020. SOLOv2: Dynamic and Fast Instance Segmentation. In *Advances in Neural Information Processing Systems*. H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 17721–17732.
- [21] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- [22] Ethan Zhang and Yi Zhang. 2009. *Average Precision*. Springer US, Boston, MA, 192–193. https://doi.org/10.1007/978-0-387-39940-9_482
- [23] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. 2017. Pyramid Scene Parsing Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2881–2890.