

BIG IMAGERY AND HIGH PERFORMANCE COMPUTING AS RESOURCES TO UNDERSTAND CHANGING ARCTIC POLYGONAL TUNDRA

C. Witharana^{1*}, M. A. E. Bhuiyan¹, A. K. Liljedahl²

¹ Dept. of Natural Resources and the Environment, University of Connecticut, Storrs, Connecticut, USA - (chandi.witharana, md.bhuiyan)@uconn.edu

² Woods Hole Research Center, Falmouth, Massachusetts, USA - aliljedahl@whrc.org

KEY WORDS: Arctic, Permafrost, Ice-wedge polygons, Deep learning, Satellite imagery, Mapping

ABSTRACT:

Permafrost thaw has been observed at several locations across the Arctic tundra in recent decades; however, the pan-Arctic extent and spatiotemporal dynamics of thaw remains poorly explained. Thaw-induced differential ground subsidence and dramatic microtopographic transitions, such as transformation of low-centered ice-wedge polygons (IWPs) into high-centered IWPs can be characterized using very high spatial resolution (VHSR) commercial satellite imagery. Arctic researchers demand for an accurate estimate of the distribution of IWPs and their status across the tundra domain. The entire Arctic has been imaged in 0.5m resolution by commercial satellite sensors; however, mapping efforts are yet limited to small scales and confined to manual or semi-automated methods. Knowledge discovery through artificial intelligence (AI), big imagery, and high performance computing (HPC) resources is just starting to be realized in Arctic science. Large-scale deployment of VHSR imagery resources requires sophisticated computational approaches to automated image interpretation coupled with efficient use of HPC resources. We are in the process of developing an automated Mapping Application for Permafrost Land Environment (MAPLE) by combining big imagery, AI, and HPC resources. The MAPLE uses deep learning (DL) convolutional neural nets (CNNs) algorithms on HPCs to automatically map IWPs from VHSR commercial satellite imagery across large geographic domains. We trained and tasked a DLCNN semantic object instance segmentation algorithm to automatically classify IWPs from VHSR satellite imagery. Overall, our findings demonstrate the robust performances of IWP mapping algorithm in diverse tundra landscapes and lay a firm foundation for its operational-level application in repeated documentation of circumpolar permafrost disturbances.

1. INTRODUCTION

Arctic permafrost - unique landscapes comprising the Earth materials that remains at or below 0°C for at least two consecutive years - occupies approximately 24% of the exposed land surface of the northern hemisphere. Ice-rich permafrost can be identified by atypical surface features called ice-wedge polygons (IWPs), which are underlain by several meter-wide and deep ice-wedges that form a network across the tundra.

Thawing of ice-rich permafrost can be seen from satellites by altered moisture and vegetation that follows the differential ground subsidence when the top of the ice-wedges thaws. The diameter of IWPs typically ranges from 5 m to 30 m and the associated microtopographic features, such as troughs or rims, are in sub-meter to ~1-meter scale (Kanevskiy et al. 2017). Vegetation and geology maps suggest that about two-thirds or more of the Arctic landscape is occupied by polygonal ground (Raynolds et al. 2019) and therefore ice-rich ground, but the exact extent and the prevailing IWP types (i.e. whether the ice wedges experienced melt or not) are largely unknown (Liljedahl et al. 2016). The microtopography associated with IWPs affects the Arctic ecosystem from local to regional scales due to the impacts on the flow and storage of water and therefore, vegetation and carbon. Over recent decades, ice-wedge degradation – transformation of low-centered polygons into high-centered polygons due to ice-wedge degradation has been documented at several locations across the Arctic tundra in the field and through localized remote sensing analyses (Steedman et al., 2017). The shift from one IWP type to the other is documented to occur in less than a decade (Liljedahl et al. 2016) with an unusual warm summer, wildfires, or human activities initiating the onset of ice-wedge degradation. Subsequent

feedback processes can either continue, amplify, dampen or even reverse the ice-wedge melt, with such processes remaining active for multiple decades. Degradation of ice wedges is a quasi-cyclic process with degradation often occurring over a shorter time scale than the formation of new permafrost (aggradation), with the latter controlled by accumulation of organic and mineral soil above the ice-wedge (Kanevskiy et al. 2017). Understanding of spatiotemporal dynamics behind the evolution of ice-wedge polygonal tundra demands for objective and detailed maps consolidating the ice wedge extent and their prevailing successional stages.

Despite the alarming signals, yet, the Arctic science community has a limited understanding of the spatiotemporal continuity of the otherwise locally observed changes. The lack of knowledge on the larger geographical extent and successional stage of IWPs introduce uncertainties to regional and pan-Arctic estimates of carbon, water, and energy fluxes. Remote sensing provides transformational opportunities to observe, monitor, and measure the Arctic polygonal landscape at multiple spatial scales and varying temporal windows. IWPs are difficult to detect in any remote sensing imagery with spatial resolution greater than 4 m. Therefore, sub-meter resolution commercial satellite imagery (e.g. DigitalGlobe, Inc.) demonstrate a greater promise in accurate delineation and characterization of ice-wedge polygonal networks. Due to IWPs' varying spectral and morphometric characteristics visual inspection and manual digitization has so far been the most widely adopted and promising method to delineate polygons from high resolution remote sensing imagery. A considerable number of local-scale studies have analyzed ice wedge degradation processes using satellite imagery, and manned-/unmanned aerial imagery/LiDAR data (Muster et al. 2013). Most of the studies

to date have relied on manual image interpretation and/or semi-automated approaches (Skurikhin et al. 2014). Therefore, there is a need and an opportunity for utilization of VHSR imagery in regional scale mapping efforts to spatio-temporally document microtopographic changes due to thawing ice-rich permafrost.

Despite the accuracy, the intense workload of manual digitization constrains the size of a study area. Sophisticated image analysis algorithms are therefore needed to automatically extract ice-wedge polygons and their prevailing stages using hundreds to thousands of commercial satellite imagery. Recently, deep learning (DL) convolutional neural nets (CNNs) has shown great potential for object instance segmentation in detecting and delineating each distinct object in an image of common objects from everyday pictures. The success of DLNNs in computer visions (CV) applications has received great interest from the remote sensing community. There has been an explosion of studies integrating DCLNN to address remote sensing classification problems spanning from general land use and land cover mapping to targeted feature extraction. Deep learning CNNs excel at object detection, semantic segmentation, and semantic object instance segmentation. A plethora of DLNN architectures have been proposed, developed, and tested. The influx of new DLNNs continues to grow. Each has its own merits and disadvantages with respect to the detection and/or classification problem at hand. DL-based approaches are being adapted to finely classify Arctic permafrost features from high-resolution remote sensing imagery (aerial (Zhang et al. 2018), satellite (Bhuiyan et al. 2019), imagery-derived digital elevation models (DEMs), and LiDAR-based DEMs.




Figure 1. (a) Circumpolar Arctic vegetation map (CAVM) showing the extent of tundra (Raynolds et al. 2019). (b) A 100 km x 100 km grid overlain on CAVM to show the geographic extent to be mapped using commercial satellite imagery, comprising tundra of Alaska, Canada, and Russia. The area under the grid covers approximately 5 million km².

The entire Arctic has been imaged at 0.5 m resolution by commercial satellite sensors (DigitalGlobe, Inc.). The image repository at the Polar Geospatial Center (PGC) provides transformational opportunities to observe, monitor, and document permafrost thaw occurring across the Arctic, which is a remote region with extremely sparse field observation network. Alaska, Canada, and Russia collectively harbor approximately 5 million km² of tundra (Figure 1). Producing a permafrost thaw map for the pan-Arctic tundra region just once would require an analysis of thousands commercial satellite image scenes. Such large-scale deployment of imagery resources requires sophisticated computational approaches to automated image interpretation coupled with efficient use of high-performance computing (HPC) resources. Besides our ongoing efforts, to the best of our knowledge, no university-led study has so far been developed a fully automated and scalable method, which is capable of accurately detecting and

characterizing surface features from sub-meter satellite imagery over large geographical areas (e.g. pan-Arctic) in an operational context.

The overarching goal of our ongoing effort is the production of the first pan-Arctic ice-wedge polygon map using large volume of commercial satellite imagery available at the Polar Geospatial Center and HPC resources. The resulting circumpolar ice-wedge polygon map will advance our understanding of the complex and interlinked processes responsible for the evolution of the pan-Arctic ice-wedge polygon tundra landscape. Here we unravel a new framework that we developed and tested – Mapping Application for Permafrost Land Environment (MAPLE) – to drive imagery-enabled Arctic permafrost science applications. We also present some of the automated ice-wedge polygon mapping results while relating to the DLNN model interoperability across different tundra vegetation types. As seen on Figure 1, we have so far completed mapping of over 10000 km² of North Slope of Alaska. The rest of the Alaska is being processed at the moment and we will expand the mapping operation to rest of tundra regions.

2. METHODS

2.1 Mapping Application for Permafrost Land Environment

A general schematics of our mapping application for permafrost land environment (MAPLE) is shown in Figure 2. While our primary target is large volumes of commercial satellite imagery, manned- and unmanned-aerial imagery can also be integrated into the workflow as needed. Our target is to derive science-ready products from imagery. Such products are not limited to ice-wedge polygons but extend to other permafrost thaw features and microtopography such as thaw slumps, lake erosion, trough, and capillaries or even vegetation (shrubs) mapping. MAPLE primarily sits on HPC resources such XSEDE (eXtreme Science and Engineering Discovery Environment) that are available for the U.S. National Science Foundation (NSF) funded research.

Accurate characterization of IWPs from VHSR imagery directly depend on the segmentation (i.e. isolation of targets from the surrounding) and classification (i.e. assigning the correct label to the targets) processes. Semantic object instance segmentation methods are designed to afford target isolation and labeling to thematic classes. Ideally, a mapping application for permafrost land environment should consist of candidate DLNN models tailored to extract different permafrost features of interest from remote sensing imagery. Among suite of target features, microtopography, thaw features, capillaries, and plant functional exhibit high priority. Given the diversity of target features and their heterogeneous characteristics coupled with semantic complexities, multiple model architectures better serve the purpose. In MAPLE, one pipeline targets on mapping ice-wedge polygons in which we utilized Mask RCNN algorithm. The pipeline is extensible and tailored to work with remote sensing imagery using high performance computing resources. This allows scalability to larger spatial extents. Mapping workflow and Mask-RCNN architecture is detailed out in Figure 2. Our mapping workflow is modular. First, input image scene is portioned into tiles of 200 pxl 200 pxl. Next, tiles are streamed to the trained model for prediction. Post processing of predictions takes place in stage 3. The predicted categorical raster is vectorized as a shapefile. All the individual shapefiles are post-processed by omitting duplicates along tile borders and

merged together to create a single shapefile corresponding to the extent of the input satellite image scene. The workflow is parallelized and optimized to work on multiple GPUs and nodes.




Figure 2. Simplified schematic of the Mapping Application for Permafrost Land Environment (MAPLE) (top). Imagery-enabled workflow embedded with the Mask R-CNN architecture (bottom).

The Mask R-CNN (He et al. 2016) serves as the key DLNN algorithm in MAPLE. It has widely been acknowledged as a promising algorithm in semantic segmentation tasks across multitude scientific domains. The Mask-RCNN is an extended method for object instance segmentation built on the Faster R-CNN with the aid of a function for predicting masks for distinct objects. The Mask R-CNN generates proposals (i.e., candidate object bounding boxes) after scanning the image, and subsequently, the model predicts the class, bounding box, and binary mask for each region of interest (ROI). The primary segments of the Mask R-CNN include: backbone architecture Residual Learning network (ResNet) for feature extraction, Feature Pyramid Network (FPN) for improving representation of objects at multiple scales, and other modules, such as Region Proposal Network (RPN) for generating (RoI), RoI Classifier for class prediction of each RoI, Bounding Box Regressor (BBR) for refining RoI, and FCN with RoIAlign and bilinear interpolation for predicting pixel-accurate mask.

2.2 Model Training

We practiced transfer learning strategy to re-train the Mask-RCNN network. Annotated data (defining and labelling regions of interest) were created using an online web tool “VGG Image Annotator” from satellite imagery comprising heterogeneous tundra types. We randomly selected 262 cropped subsets (tiles of 200 pxl by 200 pxl) (~15,000 polygons) from different tundra types (tussock, non-tussock, and sedge) considering the spectral, and spatial variability. Each file has 200 x 200 pixels. Dataset are annotated for two classes: Low-centered (LC) polygons (8962 objects) and high-centered (HC) polygons (6038 objects). The annotated tiles were randomly divided into a training dataset, validation dataset, and test dataset based on the 8:1:1 split rule. We trained the Mask R-CNN model with a mini-batch size of two image tiles, 350 steps per epoch, learning rate of 0.001, learning momentum of 0.9, and weight decay of 0.0001. To minimize overfitting, random horizontal flips augmentation was applied to introduce variability in the training data that has acceptable estimation accuracy. During calibration, the weights and biases of each neuron were estimated iteratively by

minimizing a mean squared error cost function using a gradient descent algorithm with back propagation. Training was implemented using NVIDIA V100 GPUs on XSEDE supercomputing resources. We trained the Mask R-CNN with 100 epochs. To optimize Mask R-CNN, we examined different losses, such as (a) Smooth-L1 loss, defines box regression on object detection systems, which is less sensitive to outliers, than other regression loss; (b) Mask R-CNN bounding box loss indicates the difference between predicted bounding box correction and true bounding box; (c) Mask R-CNN classifier loss estimates difference of class labels between prediction and ground truth; (d) mask binary cross-entropy loss measures (probability value between 0 and 1) the performance of a classification model by observing predicted class and actual class; (e) RPN bounding box loss identifies the regression loss of bounding boxes only when there is object and; (f) RPN anchor classifier loss indicates the difference between the predicted(RPN) and actual (closest ground truth box to the anchor box) regression.

2.3 Model Interoperability

We employed a systematic experiment to investigate the model’s immunity to landscape heterogeneity across three tundra types (tussock, non-tussock, and sedge). The experiment was based on four summer-time multi-spectral images acquired by the WorldView-2 satellite sensor (Figure 3). Pansharpened multispectral images at 0.5m were provided by the Polar Geospatial Center as orthorectified, atmospherically corrected data products. Scenes primarily comprised wet-sedge, tussock, and non-tussock tundra according to the Circumpolar Arctic Vegetation map (Raynolds et al. 2019). Candidate scenes cover 1500 km² of coastal and upland tundra (tussock, non-tussock, and sedge), from the North Slope, Alaska. We selected training study sites comprising tussock, non-tussock, sedge, and barren tundra dominant landscapes, primarily from coastal tundra region of North Slope Alaska, Canada, and Wrangler Island of Russia. The training sites provide a substantial landscape heterogeneity for model classifying and detection of ice-wedge polygons.




Figure 3. Geographic setting (left) and (b) tundra vegetation map overlain by four candidate satellite image scenes.

2.4 Accuracy Assessment

We tasked several error metrics to assess the DL model performances across tundra vegetation types.

The mean intersection over union (mIoU) between predicted and actual polygon. A mIoU score > 0.5 is considered a “good” prediction which indicates successful delineation.

$$mIoU = \frac{A_0}{A_U} \dots\dots\dots(1)$$

Absolute mean relative error (AMRE) is the mean of the relative percentage error, calculated by the normalized average

$$AMRE = \left| \frac{1}{n} \sum_{i=1}^n \left(\frac{\hat{y}_i - y_i}{y_i} \right) \right| \quad \dots\dots\dots(2)$$

Here, $\hat{y}_i$ is the number of predicted polygons, y_i is the number of actual polygons. For an unbiased model, the AMRE would be 0.

Correctness, which indicates how many of predicted positives were truly positive; Completeness, determines what percentage of actual positives were detected; F1 Score, which determines a balance between Correctness and Completeness into one value [29]. High magnitudes of Correctness suggest that there are less false positives in the classification. In addition, if the model classification always predicts positive magnitudes, Completeness will be 1, which indicates ice-wedge polygon is properly detected by the model. Moreover, an accurate prediction of ice-wedge polygon is represented by F1 score, where score of 1 specifies perfect prediction. An accurate prediction is represented by all metric values closing to 1.

$$\text{Correctness} = \frac{TP}{TP + FP} \quad \dots\dots\dots(3)$$

$$\text{Completeness} = \frac{TP}{TP + FN} \quad \dots\dots\dots(4)$$

$$\text{F1 Score} = \frac{2 * \text{Correctness} * \text{Completeness}}{\text{Correctness} + \text{Completeness}} \quad \dots\dots(5)$$

Here, true positive (TP) is the number of polygons correctly identified, false positive (FP) is the number of polygons identified by model but not true, and false negative (FN) is undetected polygons.

3. RESULTS AND DISCUSSION

We statistically evaluated the performances of the DL model in detecting and classifying IWPs. For the quantitative assessments, from each image scene, we randomly selected 40 subsets to manually delineate polygons as a reference (ground-truth polygons). The mean intersection over union (mIoU) values varied between 0.85 to 0.91 (Table 1), which indicated that predicted polygons that agree with the ground-truth polygons.

Scene	mIoU
S1	0.91
S2	0.87
S3	0.86
S44	0.85

Table 1. Summary Statistics of mean intersection over union

Zoomed-in views of the original imagery, ground truth, and model classification results show that our predicted IWPs closely matched ground-truth IWPs (Figure 4).




Figure 4. Zoomed-in views of (a) original imagery, (b) ground truth (manual delineation, blue outline) and (c) model result (yellow outline) for candidate scene 4. Imagery © [2016] DigitalGlobe, Inc.

We used three quantitative error statistics (correctness, completeness, and F1 score) to show the performances of the framework. Candidate scenes 1, 2, 3, and 4 produced high model detection accuracies for the F1 score, ranging from 0.89 to 0.96 (Figure 5, Table 2).

Scene	#of reference polygons	Correctness	Completeness	F1 Score
S1	582	0.99	89%	0.96
S2	567	1	85%	0.94
S3	579	1	83%	0.92
S4	573	1	81%	0.89

Table 2. Reported scores for correctness, completeness, and F1 measure




Figure 5. Sample views of original imagery (left) and model classification (right) for candidate scenes. Yellow outlines denote automatically detected IWPs. Imagery © [2010, 2012, 2015, 2016] DigitalGlobe, Inc.

Although all the image scenes are geographically close to each other, but they still have different tundra variations in the microtopography. Predominance of tussock sedge tundra and the high spatial resolution of imagery information provide landscape-scale variation within the original CAVM polygons throughout northern Alaska. Scene 4 (covering tundra tussock sedge) achieves mIoU 0.85 (Table 1) which still have a chance to improve model prediction by increasing more training data from that tundra region. Moreover, lake rich regions, such of Alaska's North Slope demonstrated dominant sedges tundra, which contains more detailed information which will help IWP mapping for that tundra type. Image scenes 2 and 3 represents non-tussock sedge tundra of Alaska's North Slope. Model performances for image scenes 2-3 (F1 score: 0.92-0.94) were consistent, which means training sample were sufficient to predict IWPs for non-tussock sedge tundra regions. In a similar fashion, scenes 1-4 scored high values for completeness (81-89%). In all four cases, the correctness metric scored ~ 1 , allowing less freedom for false alarms. Classification accuracies varied from 0.94 to 0.97 for candidate scenes, indicating a robust performance of the new version of the ice-wedge polygon mapping algorithm across different tundra types in northern Alaska. Results showed substantially low systematic errors (AMRE values from 0.17 to 0.23) for all candidate scenes (Figure 6). Overall, both quantitative and qualitative evaluations support the possible interoperability of the IWP mapping algorithm across different tundra assemblages in northern Alaska.




Figure 6. Absolute Mean relative Error (AMRE) for candidate scenes

Considering the geographical extent of Arctic tundra and associated landscape complexities that yet to be exposed, we make no strong claims that our transferability study is comprehensive but an exploratory effort. Arctic tundra landscapes cover spatially isolated ponds, lakes, marshes, river, and stream corridor wetlands, which representing highly heterogeneous features, varying in soil moisture, vegetation composition, elevation, surficial geology, ground ice content, soil thermal regimes and surface hydrology. Fine-scale differences in microtopography, limit the ability to comprehend local scale controls on regional to global scale patterns which, is an important factor in characterizing IWPs in arctic varying tundra areas. As Arctic tundra vegetation are spatially heterogeneous which vary significantly over small spatial scales, including additional different tundra landforms with different vegetation, hydrology, or soil characteristics could further improve the model. Further research is inevitably needed to better understand how trained models behave across tundra types. Such study can incorporate not only vegetation cover but also other factors as well, such as terrain units, soil types, hydro-climatic regimes, and permafrost characteristics. For instance, Arctic tundra types are very sensitive in summer temperature, which can cause major changes to vegetation structure via by pose spectral/textural changes in the acquired imagery. Thus, the seasonality could be an important factor deciding the appearance of ice wedge polygon on the satellite imagery because changes to spectral and textural characteristics can alter the overall semantics of the target. The model can be biased when it is trained only based on the imagery acquired in particular time window. The predictions can be suffered, if the model is given imagery from different time windows, for example early summer image vs late summer image. As much as important to understand model's interoperability across space, it is also imperative to examine how model responds to temporal variations. Operator biasness in hand-annotated data production can also negatively influence model performances. Tasking multiple operators to produce sizeable amount of quality-controlled training datasets can help improving the variability training samples and eventually leveraging the model performances.

4. CONCLUSION

We are in the process of developing a Mapping Application for Permafrost Land Environment (MAPLE) by combining big imagery, AI, and HPC resources. The MAPLE uses DLCNNs algorithms on HPCs to automatically map IWPs from VHSR commercial satellite imagery across large geographic domains. We trained and tasked a DLCNN semantic object instance segmentation algorithm to automatically classify IWPs from VHSR satellite imagery. We explored the DLCNN model interoperability across different tundra types and image scene complexities in order to understand the opportunities and challenges prior to any future circumpolar IWP mapping

applications. The DL model exhibited promising performances with high detection and classification accuracies. Also reported low systematic error for all image scenes and indicated significant improvement in model predictions across the heterogeneous tundra regions. Consideration of contextual information (e.g. edges, vegetation, shape, area, and the consistency of feature distributions) increased the reliability of the model classification and helped generalizing the DL model across tundra vegetation types. Complex topography play a vital role in controlling the spatial variation in image scenes. Effort to further refine model prediction accuracies could include a) increasing the variability of training samples with additional annotated IWPs from a larger set of tundra vegetation types, and b) explore more sophisticated image pre-processing steps such as differing data fusion approaches.

ACKNOWLEDGEMENTS

This research was support by the U.S. National Science Foundation grant #: 1722572 and 1720875. Supercomputing resources were provided by the eXtreme Science and Engineering Discovery Environment . Authors would like to thank Polar Geospatial Center, University of Minnesota for imagery support.

REFERENCES

- Bhuiyan, M.A.E., Witharana, C. and Liljedahl, A.K., 2019. Big Imagery as a Resource to Understand Patterns, Dynamics, and Vulnerability of Arctic Polygonal Tundra. *AGUFM*, pp.C13E-1374.
- He, K.; Zhang, X.; Ren, S.2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, Las Vegas, America, 27-30 Jun.
- Kanevskiy, M.; Shur, Y.; Jorgenson, T.; Brown, D.R.N.; Moskalenko, N.G.; Brown, J.; Walker, D.A.; Raynolds, M.K.; Buchhorn, M.2017. Degradation and stabilization of ice wedges: Implications for assessing risk of thermokarst in northern Alaska. *Geomorphology*, 297, 20–42.
- Liljedahl, A.K.; Boike, J.; Daanen, R.P.; Fedorov, A.N.; Frost, G.V.; Grosse, G.; Hinzman, L.D.; Iijima, Y.; Jorgenson, J.C.; Matveyeva, N.; Necsoiu, M. 2016. Pan-Arctic ice-wedge degradation in warming permafrost and its influence on tundra hydrology. *Nat. Geosci.*, 9, 312.
- Muster, S.; Heim, B.; Abnizova, A.; Boike, J. 2013. Water body distributions across scales: a remote sensing based comparison of three arctic tundra wetlands. *Remote Sens.*, 5, 1498–1523.
- Raynolds, M.K.; Walker, D.A.; Balser, A.; Bay, C.; Campbell, M.; Cherosov, M.M.; Daniëls, F.J.; Eidesen, P.B.; Ermokhina, K.A.; Frost, G.V.; Jedrzejek, B. A. 2019. raster version of the Circumpolar Arctic Vegetation Map (CAVM). *Remote Sensing of Environment*, 232,111297.2019.
- Skurikhin, A. N.; Wilson, C. J.; Liljedahl, A.; Rowland, J. C. 2014. Recursive active contours for hierarchical segmentation of wetlands in high-resolution satellite imagery of arctic landscapes. *Southwest Symposium on Image Analysis and Interpretation*, 137–140.

Steedman, A. E.; Lantz, T. C.; Kokelj, S. V. 2017. Spatio-temporal variation in high-centre polygons and ice-wedge melt ponds, Tuktoyaktuk coastlands, Northwest Territories. *Permafr. Periglac. Process.*, 28, 66–78.

Zhang, W.; Witharana, C.; Liljedahl, A.K.; Kanevskiy, M. 2018. Deep Convolutional Neural Networks for Automated Characterization of Arctic Ice-Wedge Polygons in Very High Spatial Resolution Aerial Imagery. *Remote Sens.*, 10, 1487.