Brief communication: Rapid machine learning-based extraction and measurement of ice wedge polygons in airborne lidar data

Charles J. Abolt[1,2], Michael H. Young[2], Adam A. Atchley[3], Cathy J. Wilson[3]

[1]Department of Geological Sciences, The University of Texas at Austin, Austin, TX, USA
5 [2]Bureau of Economic Geology, The University of Texas at Austin, Austin, TX USA
[3]Earth and Environmental Sciences Division, Los Alamos National Laboratory, Los Alamos, NM, USA

Correspondence to: Charles J. Abolt (chuck.abolt@beg.utexas.edu)

Abstract. We present a workflow for rapid delineation and microtopographic characterization of ice wedge polygons within
10 high-resolution digital elevation models. The workflow, which is extensible to other forms of remotely sensed imagery,
incorporates a convolutional neural network to detect pixels representing troughs. A watershed transformation is then used to
segment imagery into discrete polygons. Regions of non-polygonal terrain are partitioned out using a simple post-processing
procedure. Results from training and validation sites at Barrow and Prudhoe Bay, Alaska demonstrate robust performance in
diverse tundra landscapes. The methodology permits fast, spatially extensive measurements of polygonal microtopography
15 and trough network geometry.

1 Introduction and Background

The objective of this research is to develop and report on a workflow for rapid delineation and microtopographic
analysis of ice wedge polygons in airborne lidar data. Ice wedge polygons are the surface expression of ice wedges, a form of
ground ice nearly ubiquitous to coastal tundra environments in North America and Eurasia (Leffingwell, 1915; Lachenbruch,
20 1962). Several thousand ice wedge polygons of typical size may occupy a single square kilometer of terrain, motivating our
development of an automated approach to map polygonal boundaries. The key innovation in our method is the use of a
convolutional neural network (CNN), a variety of machine learning algorithm, to identify troughs (*i.e.*, polygonal boundaries)
within lidar-derived imagery. The same techniques could be applied to any form of remotely sensed data with sufficient spatial
resolution; however, using a digital elevation model (DEM) as input, our approach permits the extraction of microtopographic
25 attributes from entire populations of ice wedge polygons at the kilometer scale or greater.

High-resolution inventories of polygon geomorphology are of hydrologic and ecologic interest because decimeter-
scale variability in polygon microtopography can drive pronounced changes to soil drainage (Liljedahl et al., 2016) and surface
emissions of CO_2 and CH_4 (Lara et al., 2015; Wainwright et al., 2015). Several previous geospatial analyses of polygon
geomorphology have been motivated by a recent acceleration in the rates at which low-centered polygons (LCPs) are converted
30 into high-centered polygons (HCPs), a process driven by permafrost degradation which improves soil aeration and stimulates

enhanced emissions of CO_2 (Jorgenson et al., 2006; Raynolds et al., 2014; Jorgenson et al., 2015; Liljedahl et al., 2016). In these prior studies, the most common form of imagery used to demonstrate this transformation has been historic aerial photography. Thus far, analysis of available datasets has demonstrated a consistent uptick, beginning around 1989, in landscape-scale HCP formation throughout the high Arctic; however, precise rates of geomorphic change have been difficult

5 to quantify, as the surveys primarily have relied on proxy indicators, such as the presence of water in deepening HCP troughs (Steedman et al., 2017). In a related effort to characterize contemporary polygon microtopography, a landcover map of LCP and HCP occurrence across the Arctic coastal plain of northern Alaska was recently developed using multispectral imagery from the Landsat8 satellite at 30 m resolution (Lara et al., 2018). This dataset offers a static estimate of variation in polygonal form over unprecedented spatial scales; however, geomorphology was inferred from the characteristics of pixels larger than

10 typical polygons, preventing inspection of individual features.

Higher-resolution approaches to explicitly segment imagery into discrete polygons have typically been motivated by efforts to analyze trough network geometry. On both Earth and Mars, for example, paleo-environmental conditions in remnant polygonal landscapes have been inferred by comparing parameters such as trough spacing and orientation with systems in modern periglacial terrain (*e.g.*, Levy et al., 2009; Ulrich et al., 2011; Ewertowski et al., 2017). An early semi-automated

15 approach to delineating Martian polygons from satellite imagery was developed by Pina et al. (2006), who first used morphological image processing operations to emphasize polygonal troughs, then applied a watershed transformation (discussed in Section 4.1.3) to identify discrete polygons. This workflow was later applied to lidar-derived imagery from a landscape outside Barrow, AK by Wainwright et al. (2015), but in their experience and our own, robust results at spatial scales approaching a square kilometer or greater were elusive.

20 Our application of CNNs to the task of identifying polygonal troughs was inspired by the remarkable solutions that CNNs recently have permitted to previously intractable image processing problems. As the performance of graphics processing units (GPUs) has steadily improved over the last decade, CNNs have demonstrated unprecedented skill at tasks analogous to trough delineation, such as cell membrane identification in biomedical images (Ciresan et al., 2012) or road extraction from satellite imagery (Kestur et al., 2018). In training a CNN to analyze lidar imagery from polygonal terrain, our goal was to

25 improve the speed and accuracy of automated polygon delineation in modern tundra environments. Because our method extracts polygons from a high resolution DEM, in enables thorough, landscape-scale inventories of polygonal microtopography, and we anticipate that it will permit precise monitoring of contemporary surface deformation in landscapes covered by repeat airborne surveys.

2 Study areas and data acquisition

30 To demonstrate the flexibility of our approach, we applied it simultaneously at study sites near Barrow and Prudhoe Bay, Alaska, two setting with highly divergent ice wedge polygon geomorphology, ~300 km distant from one another.

2.1 Barrow

The first study site is located within 10 km of the Beaufort Sea coast (Fig S1A) in the Barrow Environmental Observatory, operated by the National Environmental Observatory Network (NEON). Mean elevation is less than 5 m above sea level, and vegetation consists of uniformly low-growing grasses and sedges. Mesoscale topography is mostly flat but

5   marked by depressions up to 2 m deep associated with draws and drained lake beds (Fig S1B). Microtopography at the site reflects nearly ubiquitous ice wedge development, which becomes occluded in some of the depressions. Ice wedge polygons are of complex geometry and highly variable area, ranging from ~10 m^2 to >2000 m^2. An airborne lidar survey was flown in August 2012 as part of the U.S. Department of Energy's Next Generation Ecosystems Experiment-Arctic program (https://ngee-arctic.ornl.gov/). The resulting point cloud was processed into a 50 cm horizontal resolution digital elevation

10   model (DEM) with an estimated vertical accuracy of 0.145 m (Wilson et al., 2012).

2.2 Prudhoe Bay

The second site is approximately 300 km east of the first and farther inland, located ~40 km south of Prudhoe Bay, AK (Fig S1A). As at Barrow, vegetation consists almost exclusively of low and even-growing grasses and sedges. Mesoscale topography is extremely flat, with a slight (<4%) dip toward the northwest (Fig S1C). Ice wedge troughs are apparent

15   throughout the training site, demarcating polygons of more consistent area than those of Barrow, ~400-800 m^2. Airborne lidar data was acquired in August 2012 by the Bureau of Economic Geology at the University of Texas at Austin and subsequently processed into a 50 cm resolution DEM (Paine et al., 2015). Vertical accuracy was estimated at 0.10 m.

3 Methods

3.1 Polygon delineation algorithm

20   A high-level summary of our workflow is presented in Fig 1, using a sample of the data from Barrow. In the first (pre-processing) stage, regional trends were removed from a DEM (Fig S2), generating an image of polygonal microtopography (Fig 1A). Next, the microtopographic information was processed by a CNN, which was trained to use the 27×27 neighborhood surrounding each pixel to assign a label of "trough" or "not trough" (Fig 1B). Each non-trough pixel was then assigned a negative intensity proportional to its Euclidean distance from the closest trough, generating a grayscale image

25   analogous to a DEM of isolated basins, in which the polygonal troughs appear as ridges (Fig 1C). Subsequently, a watershed transform was applied to segment the image into discrete ice wedge polygons (Fig 1D). These steps, and a post-processing algorithm used to remove non-polygonal terrain from the final image, are described in detail below.

3.1.1 Pre-processing

In the pre-processing stage, regional topographic trends were estimated by processing the DEM with a 2D filter, which assigned to each pixel the mean elevation within a 20 m (40 pixel) radius. Microtopography was estimated by subtracting this regional topography from the DEM (Fig S2). In preparation for passing the data to the CNN, microtopography was then
5   converted to 8-bit gray scale imagery. The minimum intensity (0) was assigned to depressions of 0.7 m or greater, and the maximum intensity (255) was assigned to ridges of 0.7 m or greater.

3.1.2 Convolutional neural network

The function of the CNN in our workflow was to identify pixels likely to represent troughs. Conceptually, a CNN is a classification tool which accepts images of a fixed size (in our case, 27×27 grayscale arrays) as input and generates categorical
10   labels as output. The CNN determines decision criteria through training with a set of manually labeled images. The architecture of a CNN consists of a user-defined sequence of components, or layers, which take inspiration from the neural connections of the visual cortex. Our CNN was purposefully constructed with an architecture of minimal complexity, to maximize the efficiency of training and application. Here we briefly summarize the function of each layer in our CNN; for more detailed description, the reader is directed to Ciresan et al. (2012).

15   The most important components of our CNN were a single convolutional layer, a max-pooling layer, and two fully connected layers (Table 1). In the convolutional layer, a set of 2D filters was applied to the input image, generating intermediate images in which features including concavities, convexities, or linear edges were detected. The max-pooling layer downsized these intermediate images from 27×27 to 9×9 arrays by selecting the highest intensity pixel in a moving 3×3 window with a stride of 3 pixels. Each pixel in the downsized intermediate images was then passed as an input signal to the fully connected
20   layers, which functioned identically to standard neural networks. Two additional components of our CNN were Rectified Linear Unit (ReLU) layers, which enhance non-linearity by reassigning a value of zero to any negative signals output by a preceding layer, and a softmax layer, which converted the output from the final fully connected layer into a probability for each categorical label (*i.e.*, trough or not trough).

During training, the weights of the 2D filters in the convolutional layer and the activation functions of the neurons in
25   the fully connected layers were optimized to correctly predict the labels in a training deck of images. Our workflow was designed to generate the training deck by processing square tiles of manually-labeled imagery with 100 m (200 pixel) edges. In each of these tiles, trough pixels were delineated by hand in a standard raster graphics editor, a process that required ~1 hour per tile (Fig S3A). Our algorithm imported these tiles, identified the geographic coordinates of each pixel identified as a trough, then extracted the 27×27 neighborhood surrounding each trough pixel from the 8-bit microtopographic imagery. This
30   procedure generated several thousand thumbnail-sized images centered on a trough from each manually delineated tile. Subsequently, an equal number of pixels not labeled as troughs were selected at random, and the image extraction procedure was repeated. Just prior to training, 25% of the training deck assembled by these methods was set aside to be used for validation.

Due to the compact architecture of our CNN, training and execution speeds were rapid. Processing information from ~30,000 images, classification accuracies >99% on the training deck and >98% on the validation deck were obtained in less than ten minutes on a single GPU. Application of the trained CNN to a square kilometer of gridded lidar data required ~90 seconds. This speed permitted us to iteratively train the CNN, identify regions of the survey area where it performed poorly, and supplement the training data with labeled imagery from those regions. Optionally, our workflow at this stage also permitted supplementation of the training deck with individual instances of problematic features, such as LCP centers or especially narrow troughs (Fig S3B). This permitted us to increase the skill of the CNN in a targeted manner without the time requirements of manually labeling a full 100 m tile of imagery.

3.1.3 Polygon extraction

After applying the CNN to classify all pixels at a site as trough or not trough, we extracted discrete ice wedge polygons through application of several standard image processing operations. The first step was elimination of "salt and pepper" noise in the binary image, which we accomplished by eliminating all contiguous sets of trough-identified pixels with an area < 20 m^2. Next, we assigned to every non-trough pixel a negative intensity equal to its Euclidean distance from the nearest trough. This created an intermediate image in which each ice wedge polygon appeared as a valley, surrounded on all or most sides by ridges representing the trough network (Fig 1C). At this stage, to prevent over-segmentation, valleys with maximum depths of three pixel-widths or fewer were then identified and merged with the closest neighbor through morphological reconstruction (Soille, 1992). The effect of this procedure was to ensure that the algorithm would not delineate any polygon whose center did not include a point greater than 1.5 m from the troughs. Next, watershed segmentation was then applied to divide the valleys into discrete polygons. Our use of this operation was inspired by its incorporation in the polygon delineation method developed by Pina et al. (2006). Conceptually, this procedure was analogous to identifying the up-gradient region or area of attraction surrounding each local minimum.

3.1.1 Post-processing

In the final stage of delineation, we partitioned out regions of a survey area that had been segmented using techniques described above but were unlikely to represent true ice wedge polygons. For example, polygons were eliminated from the draw in the southern half of the Barrow training site (Fig S1B), where microtopography is too occluded to permit accurate delineation. Toward this aim, our algorithm tabulated within the boundaries of each delineated polygon (black lines in Fig 1D) the number of pixels that had been identified by the CNN as troughs (white pixels in Fig 1B). It then eliminated contiguous regions greater than 10,000 m^2 (representing the area of several large polygons) where less than half the boundary pixels of the contained polygons had been labeled as troughs. This procedure had the strengths of being conceptually simple and providing a deterministic means of partitioning non-polygonal terrain from the rest of a survey area.

3.2 Microtopographic analysis

To demonstrate our workflow's capabilities for microtopographic analysis, we developed a simple method for measuring the relative elevation at the center each delineated polygon. This served as a rough proxy for LCP or HCP form. In each polygon, we first calculated the Euclidean distance from the trough network of all interior pixels, then identified the pixel

5   most distal from a trough. Elevation was extracted from this pixel and compared with mean elevation in the troughs surrounding the polygon. We thus assigned a number to each polygon which consistently was lower in LCPs and higher in HCPs.

3.3 Case study experimental design

We conducted our case study in two phases: calibration and validation. Calibration was performed at two 1 km^2 sites,

10 sampled from the Barrow and Prudhoe Bay DEMs. Leveraging the rapid training and application times of our CNN, we took an iterative approach to extracting polygons from both sites: after training and applying an initial model, new training data were introduced from regions of poor performance to improve skill. After several iterations, our final CNN was trained using manually labeled data from three 100 m tiles at the Barrow training site and one at the Prudhoe Bay site, representing 3% and 1% of the training plots, respectively. The non-trough training data were also supplemented with some additional examples of

15 LCP centers from Barrow, to discourage inaccurate labeling as troughs (Fig S3B). After training the CNN, we tested its extensibility by applying it at two validation sites, also 1 km^2 in area, selected from the same landscapes. Once delineation was complete, to illustrate capabilities for microtopographic analysis, we calculated the relative elevations of polygon centers at the Prudhoe Bay training site, which was chosen for the tendency of its polygons to cluster into areas with high concentrations of either LCPs or HCPs.

20 ### 4 Results and discussion

Results from the Barrow training site demonstrate the capability of our workflow to extract polygons of highly variable geometry (Fig 2A). The results also demonstrate success of our post-processing workflow in partitioning out a streambed traversing the southern half of the site and a lakebed in the northeast quadrant. The final delineation makes immediately apparent the complex distribution of polygon type at the kilometer scale. For example, just above the draw in the

25 southwest quadrant is a cluster of comparatively small HCPs that stands in stark contrast to the larger LCPs and intermediate polygons dispersed elsewhere. The complexity of the delineated network underscores the value of an automated approach to polygon extraction, as the manual processing necessary to achieve similar results would require prohibitively large amounts of time. Moreover, the site contains several areas in which larger polygons typically are subdivided by secondary and tertiary ice wedges. Although polygonal boundaries are somewhat ambiguous in such situations, our algorithm provides a consistent,

30 deterministic method of interpretation, which is set by the user's delineation of troughs in the training dataset.

Results from the Prudhoe Bay training site, using the same CNN, demonstrate the ease with which our workflow can

be extended across diverse landscapes. We note that, while we had initially trained our CNN using three tiles of manually labeled data from Barrow, successful extension to Prudhoe Bay required incorporation of only one additional tile, from a region in the southeast quadrant characterized by LCPs with exceptionally narrow troughs. The final delineation (Fig 2B) emphasizes the far more consistent size and geometry of polygons relative to Barrow, which may be associated with the simpler

5   mesoscale topography (Lachenbruch, 1962). The ice wedge network at Prudhoe Bay appears to be less riddled with secondary ice wedges than at Barrow, but some examples are evident in HCPs.

       Results from the validation sites demonstrate similar skill in polygon delineation without incorporation of any new training data. At the Barrow validation site (Fig S4A), the algorithm partitioned out another stream bed and three regions where surface inundation prevented accurate delineation of polygonal microtopography. Elsewhere, a mixture of LCPs, HCPs, and

10   intermediate polygons reminiscent of the training site was successfully delineated. At the Prudhoe Bay validation site (Fig S4B), the algorithm partitioned out a thaw lake and the flood plain of a braided stream, in which no ice wedge troughs are present. It also partitioned out a zone just north of the lake where ice wedges may be present in the subsurface, but microtopography is difficult to distinguish. Through the remainder of the area, the algorithm delineated a set of LCPs and HCPs of similar size and geometry to the training site, but with somewhat subtler troughs.

15   The capacity of our workflow for high-resolution microtopographic analysis is demonstrated in estimates of relief at the polygon centers at the Prudhoe Bay training site (Fig 3A). Through relatively simple calculations, we highlight the presence of a distinct cluster of LCPs in the southeast quadrant, which emerges from the mixture of HCPs and intermediate polygons present through the remainder of the site. Using this approach, our procedure provides a means to estimate polygonal geomorphology with unprecedented detail at the kilometer scale, which may improve predictions of soil wetness and polygon-

20   scale fluxes of CO_2 and CH_4 (Lara et al., 2015; Wainwright et al., 2015). Simultaneously, the delineation performed at the same site provides a means to analyze trough geometric parameters, such as spacing and orientation. We acknowledge that, because the surface expression of ice wedges is sometimes subtle or non-existent, ground-based methods remain the highest-accuracy approach to mapping ice wedge networks (Lousada, 2018). Nonetheless, by segmenting machine-delineated polygonal boundaries into individual troughs (Fig 3B), very large datasets can rapidly be generated to estimate statistics at

25   spatial scales unattainable through on-site surveying.

5 Conclusions

      Convolutional neural networks enable rapid delineation of ice wedge networks at unprecedented spatial scales using airborne lidar datasets. Successful application of CNNs is facilitated by employing relatively simple neural architecture, which permits rapid training, testing, and incorporation of new data to improve skill. A single CNN is capable of extracting polygons

30   of highly variable area and geometry from diverse tundra landscapes, using a training workflow that can be completed in a single afternoon. Potential applications for the technology include: generation of high-resolution maps of land cover by polygon type, precise quantification of microtopographic deformation in areas covered by repeat airborne surveys, and rapid

extraction of trough parameters including spacing and orientation. These capabilities can improve understanding of environmental influences on network geometry and facilitate assessments of contemporary landscape evolution in the Arctic.

Acknowledgements

We are grateful for the support provided for this research, which included: the Next Generation Ecosytem Experiments Arctic (NGEE-Arctic) project (DOE ERKP757) funded by the Office of Biological and Environmental Research in the US Deparment of Energy Office of Science, and the NASA Earth and Space Science Fellowship program, for an award to CJA. We thank Dylan Harp (Los Alamos National Laboratory) for constructive conversations during manuscript preparation. We acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper.

Data availability

Data and code are available through the corresponding author. Upon publication of the final paper, data and code will be uploaded to a Zenodo repository.

References

Ciresan, D., Giusti, A., Gambardella, L. M. and Schmidhuber, J.: Deep neural networks segment neuronal membranes in electron microscopy images, in: Advances in Neural Information Processing Systems 25, Pereira, F., Burges, C. J. C., Bottou, L. and Weinberger, Q (eds.), Curran Associates, Inc., 2843-2851, 2012.

Ewertowski, M. W., Kijowski, A., Szuman, I., Tomczyk, A. M. and Kasprzak, L.: Low-altitude remote sening and GIS-based analysis of cropmarks: Classification of past thermal-contraction crack polygons in central western Poland, Geomorphology, 293, 418-432, 2017.

Jorgenson, M. T., Shur, Y. L. and Pullman, E. R.: Abrupt increase in permafrost degradation in Arctic Alaska, Geophysical Research Letters, 33(2), doi:10.1029/2005GL024960, 2006.

Jorgenson, M. T., Kanevskiy, M., Shur, Y., Moskalenko, N., Brown, D. R. N., Wickland, K., Striegl, R. and Koch, J.: Role of ground ice dynamics and ecological feedbacks in recent ice wedge degradation and stabilization, Journal of Geophysical Research: Earth Surface, 120(11), 2280–2297, doi:10.1002/2015JF003602, 2015.

Kestur, R., Farooq, S., Abdal, R., Mehraj, E., Narasipura, O. and Mudigere, M.: UFCN: A fully convolutional neural network for road extraction in RGB imagery acquired by remote sensing from an unmanned aerial vehicle, Journal of Applied Remote Sensing, 12, 016020, doi:10.1117/1.JRS,12.016020, 2018.

Lachenbruch, A. H.: Mechanics of thermal contraction cracks and ice-wedge polygons in permafrost, Special Paper, Geological Society of America, New York, 1962.

Lara, M. J., McGuire, A. D., Euskirchen, E. S., Tweedie, C. E., Hinkel, K. M. and Skurikhin, A. N.: Polygonal tundra geomorphological change in response to warming alters future CO_2 and CH_4 flux on the Barrow Peninsula, Global Change Biology, 21, 1634-1651, 2015.

Lara, M. J., Nitze, I., Grosse, G. and McGuire, A. D.: Tundra landform and vegetation trend maps for the Arctic Coastal Plain of northern Alaska, Scientific Data, 5, 180058, 2018.

Leffingwell, E. K.: Ground-ice wedges: The dominant form of ground-ice on the north coast of Alaska, J. Geol., 23, 635-654, 1915.

Levy, J. S., Marchant, D. R. and Head, J. W.: Thermal contraction crack polygons on Mars: A synthesis from HiRISE, Phoenix, and terrestrial analogue studies, Icarus, 226, 229-252, 2010.

Liljedahl, A. K., Boike, J., Daanen, R. P., Fedorov, A. N., Frost, G. V., Grosse, G., Hinzman, L. D., Iijma, Y., Jorgenson, J. C., Matveyeva, N., Necsoiu, M., Raynolds, M. K., Romanovsky, V. E., Schulla, J., Tape, K. D., Walker, D. A., Wilson, C. J., Yabuki, H., and Zona, D.: Pan-Arctic ice-wedge degradation in warming permafrost and its influence on tundra hydrology, Nat. Geosci., 9, 312–318, https://doi.org/10.1038/ngeo2674, 2016.

Lousada, M., Pina, M., Vieira, G., Bandeira, L. and Mora, C.: Evaluation of the use of very high resolution aerial imagery for accurate ice-wedge polygon mapping (Adventdalen, Svalbard), Science of the Total Environment, 615, 1574-1583, 2018.

Paine, J. G., Andrews, J. R., Saylam, K. and Tremblay, T. A.: Airborne LiDAR-based wetland and permafrost-feature mapping on an Arctic Coastal Plain, North Slope, Alaska, in: Remote Sensing of Wetlands: Applications and Advances, Tiner, R. W., Klemas, V. V., Lang, M. W. (eds.), CRC Press, Boca Raton, FL, USA, 413-434, 2015.

Pina, P., Saraiva, J., Bandeira, L. and Barata, T.: Identification of Martian polygonal patterns using the dynamics of watershed contours, in: Image Analysis on Recognition, Campilho, A., Kamel, M. (eds.), Springer-Verlag, Berlin-Heidelberg, 691-699, 2006.

Raynolds, M. K., Walker, D. A., Ambrosius, K. J., Brown, J., Everett, K. R., Kanevskiy, M., Kofinas, G. P., Romanovsky, V. E., Shur, Y. and Webber, P. J.: Cumulative geoecological effects of 62 years of infrastructure and climate change in ice-rich permafrost landscapes, Prudhoe Bay Oilfield, Alaska, Global Change Biology, 20(4), 1211–1224, doi:10.1111/gcb.12500, 2014.

Soille, P.: Morphological Image Analysis, Springer-Verlag: Berlin Heidelberg, 2004.

Steedman, A. E., Lantz, T. C. and Kokelj, S. V.: Spatio-temporal variation in high-centre polygons and ice-wedge melt ponds, Tuktoyaktuk coastlands, Northwest Territories, Permafrost and Periglacial Processes, 28, 66-78, 2017.

Ulrich, M., Hauber, E., Herrzschuh, U., Hartel, S. and Schirrmeister, L.: Polygon pattern geomorphometry on Svalbard (Norway) and western Utopia Planitia (Mars) using high resolution stereo remote-sensing data, Geomorphology, 134, 197-216, 2011.

Wainwright, H. M., Dafflon, B., Smith, L. J., Hahn, M. S., Curtis, J. B., Wu, Y., Ulrich, C., Peterson, J. E., Torn, M. S. and Hubbard, S. S.: Identifying multiscale zonation and assessing the relative importance of polygon geomorphology on carbon fluxes in an Arctic tundra ecosystem, Journal of Geophysical Research: Biogeosciences, 120(4), 788–808, doi:10.1002/2014JG002799, 2015.

5   Wilson, C., Gangodagamage, C. and Rowland, J.: Digital elevation model, 0.5 m, Barrow Environmental Observatory, Alaska, 2012, Next Generation Ecosystem Experiments Arctic Data Collection, Oak Ridge National Laboratory, US Department of Energy, doi:10.5440/1109234, 2013.

Table 1.   Architecture of our CNN.

Layer	Type	Neurons
1	Convolutional	8 arrays of 27×27
2	ReLU†	8 arrays of 27×27
3	Max-pooling	8 arrays of 9×9
4	ReLU	8 arrays of 9×9
5	Fully-connected	64
6	ReLU	64
7	Fully-connected	2
8	ReLU	2
9	Softmax	2

† - ReLu – rectified linear unit
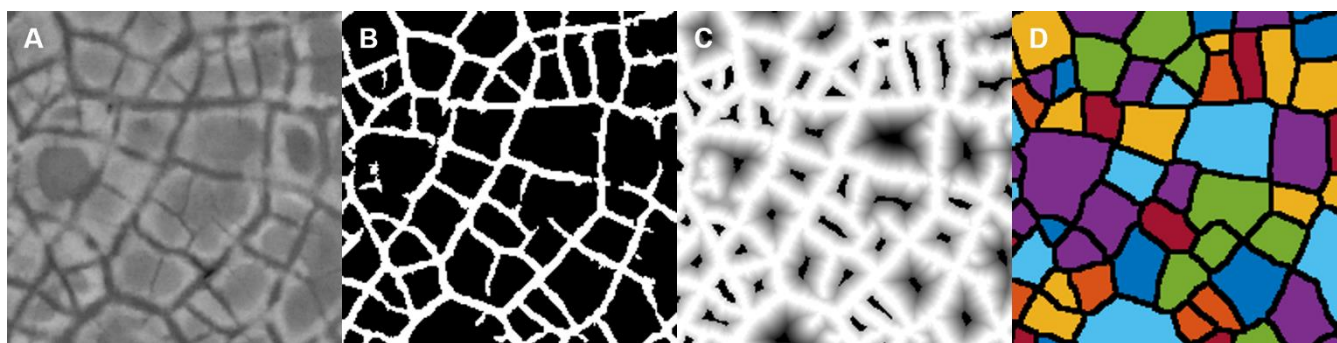
Figure 1. High-level schematic of our workflow. The input to the CNN is 8-bit grayscale imagery representing microtopography (A), estimated by removing regional trends from the lidar DEM (Fig S2). The CNN identifies pixels likely to represent troughs (B). Each non-trough pixel is assigned a negative intensity equal to its distance from a trough (C) and a watershed transformation is applied to segment the image into discrete polygons (D) (colors randomly applied to emphasize polygonal boundaries).
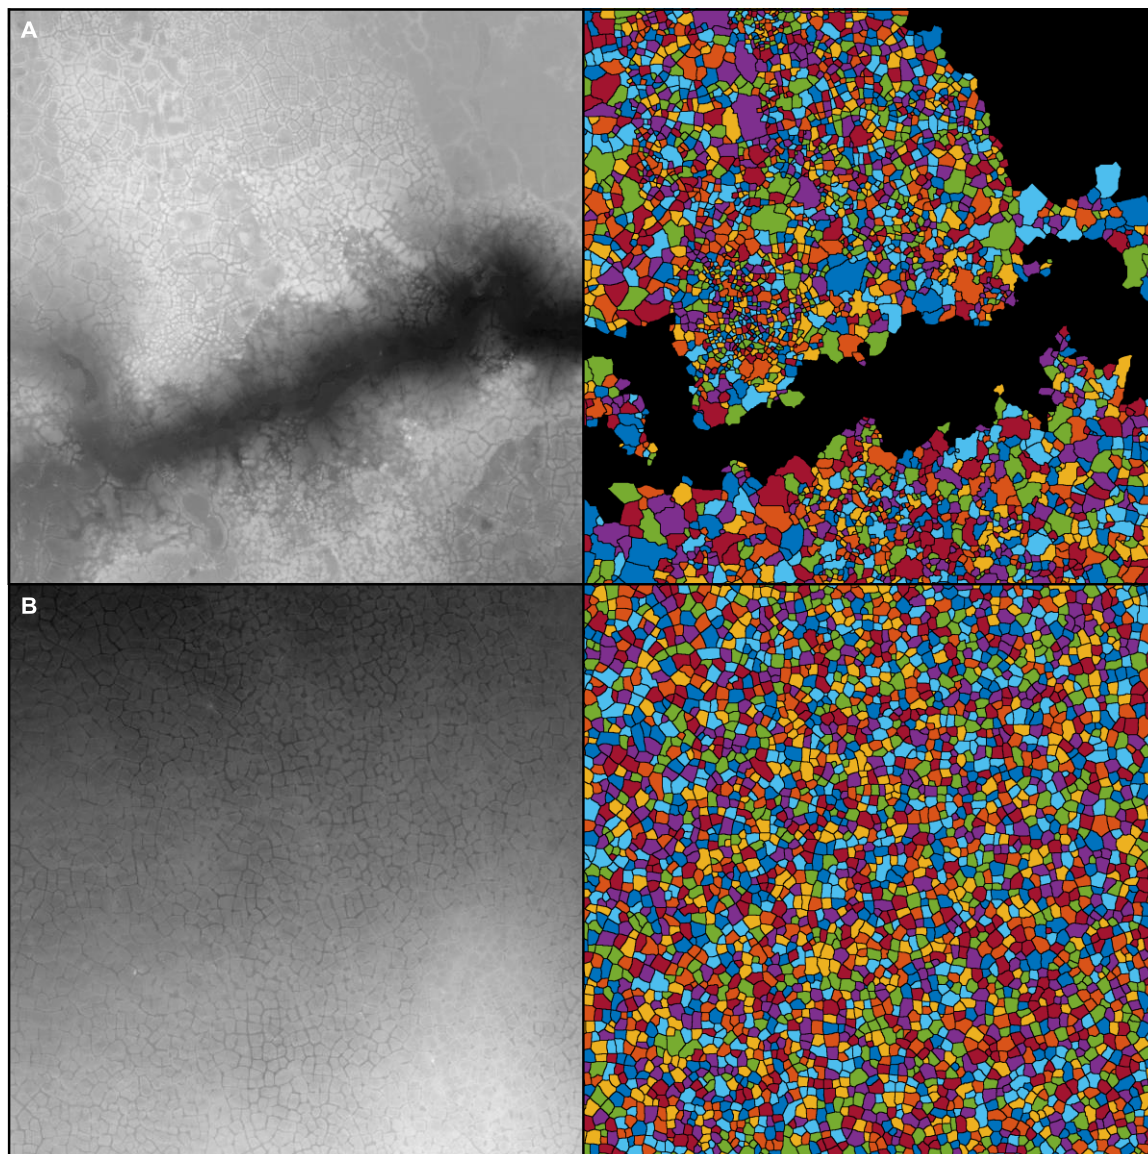
Figure 2. Delineation of ice wedge polygons using a single CNN at Barrow (A) and Prudhoe Bay (B) training sites. DEMs are shown on the left, and polygons are randomly colorized on the right.
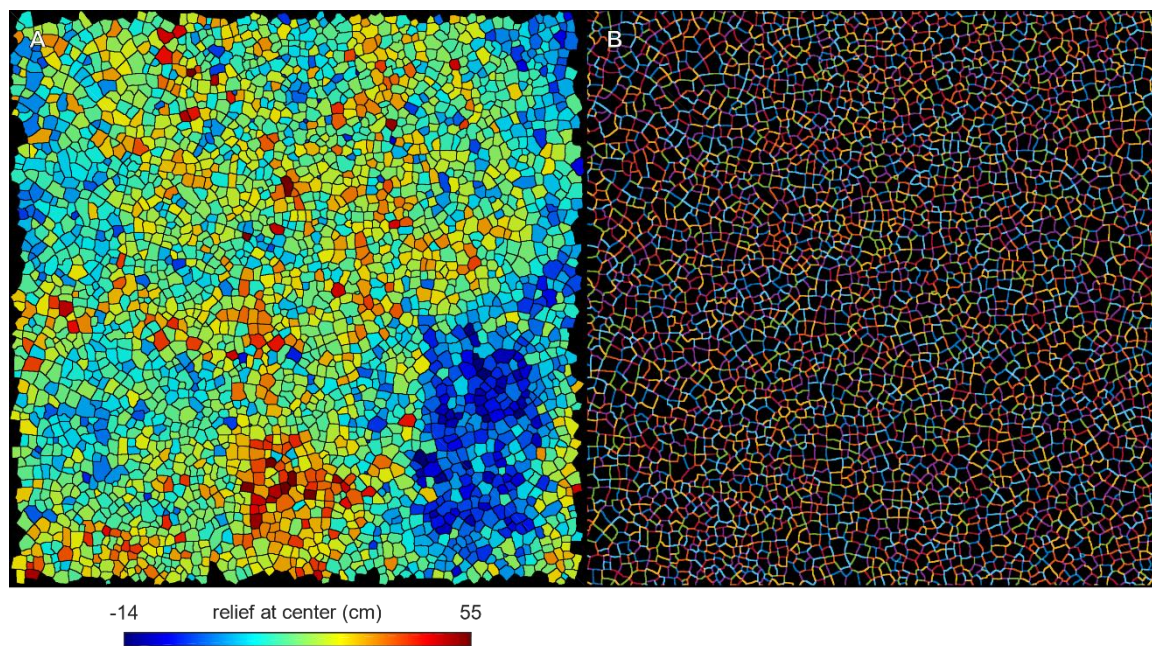
Figure 3.   Relative elevations of polygon centers at Prudhoe Bay training site (A). Automated segmentation of the primary trough network at the Prudhoe Bay training site, randomly colorized (B).