Automated Hailpad Dent Detection and Segmentation Using Machine Learning

Western Canadian Severe Storms Laboratory

Kevin Manka, Daniel Butt, Connell Miller, Julian Brimelow

Background

Hailpad networks are widely used as a simple and low-cost approach to in-situ hailstone measurement by capturing indentations made during severe convective storms.[1] However, the subsequent process of analyzing these indentations is often time-intensive and subjective.

Example of a mounted hailpad.

Objectives

- Establish a novel automated hailpad analysis pipeline that employs image processing and machine learning techniques to threshold, segment, and measure the axes and maximum depths of dents from 3D hailpad scans with limited manual intervention.
- Reduce the time required to analyze a hailpad from hours to minutes.

Data

Hailpads are scanned using a Creaform MetraSCAN BLACK+ Elite 3D handheld laser scanner and a C-Track optical tracker.

Access to real, labelled hailpads is a considerable limitation that arises in constructing a large and varied dataset required for a performant machine learning model. As such, the datasets are instead composed of synthetic hailpad binary masks that can easily be generated and labelled en masse.

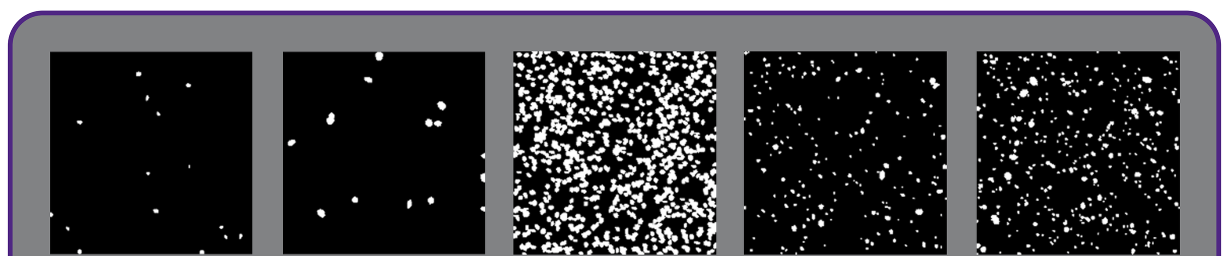

For training and validation sets, 5 categories of synthetic hailpads encompass the varying hailpad dent distributions observed in Northern Hail Project field deployments. They vary along the following parameters:

- **Diameter range**.
- **Axis variation** (v). The maximum possible reduction in the minor axis a relative to the major axis b, where $a = b(1 - vx)$, $x \sim U(0, 1)$.
- **Distribution** (uniform or exponential).
- **Dent**, **sample**, and **hailpad counts**.

Each generated dent ellipse is distorted using random noise and drawn to the main binary mask in addition to its own mask, unless it is fully enclosed by another dent. In such a case, it is removed entirely to avoid generating misleading training samples that suggest a separate dent exists within a single continuous dent.

1 POLYSTYRENE FOAM HAILPAD

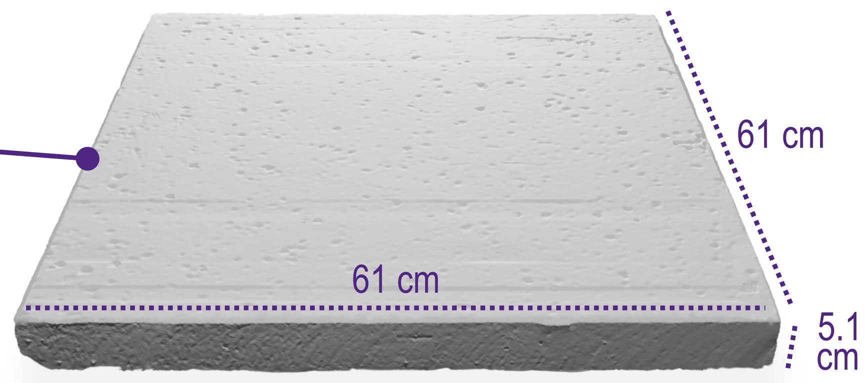

2 DEPTH MAP OF SURFACE LASER SCAN

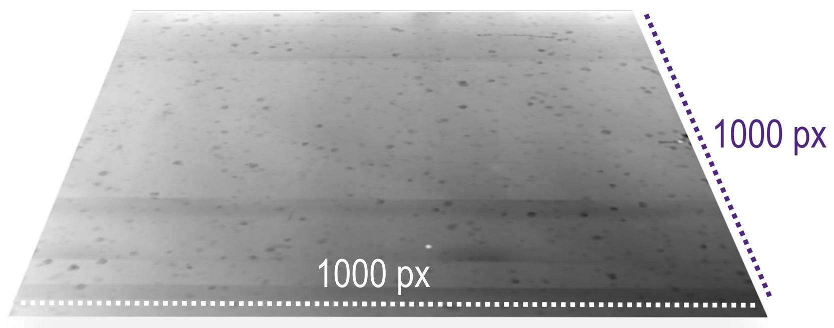

3 CLAHE[2] (ACCENTUATES LOCAL CONTRAST OF SHALLOWER DENTS)

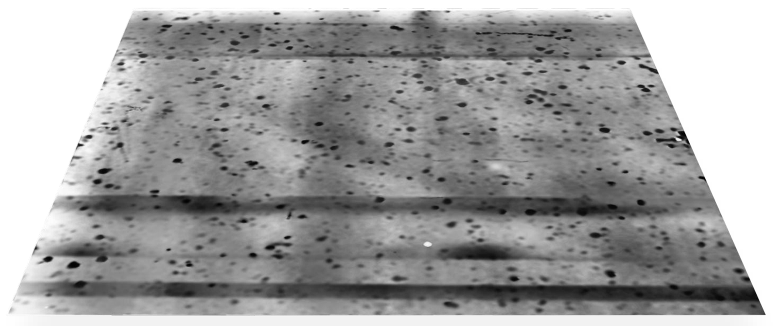

4 ADAPTIVE THRESHOLDING (SEPARATES DENTS FROM BACKGROUND)

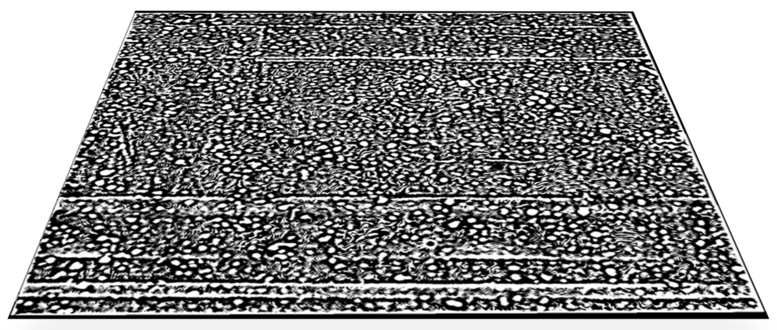

5 EROSION-DILATION AND AREA FILTERING (REMOVE NOISE)

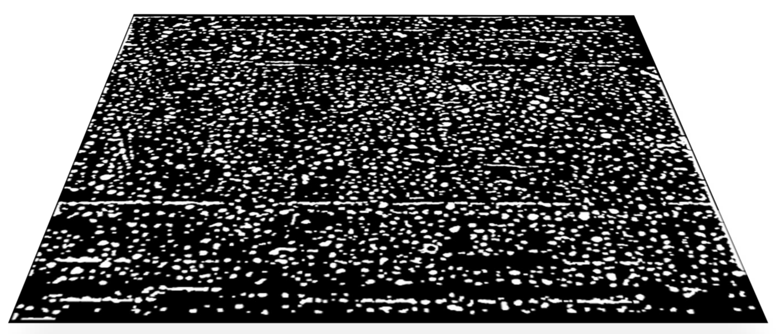

6 POINT QUERY SELECTION

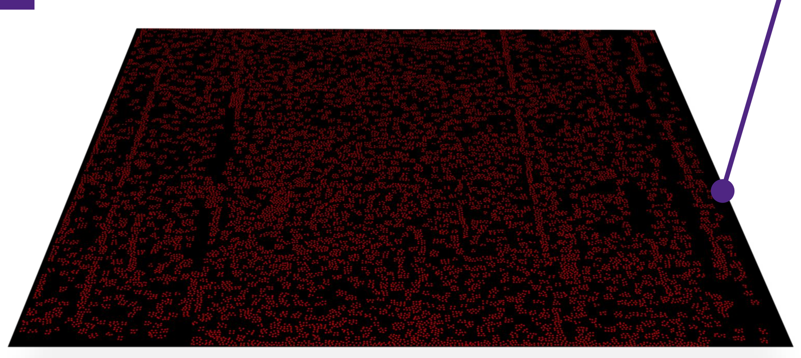

7 REFINED DENT PREDICTIONS FROM POINT QUERIES

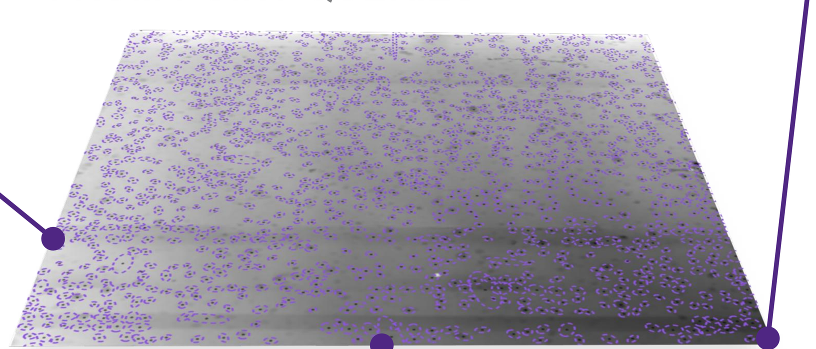

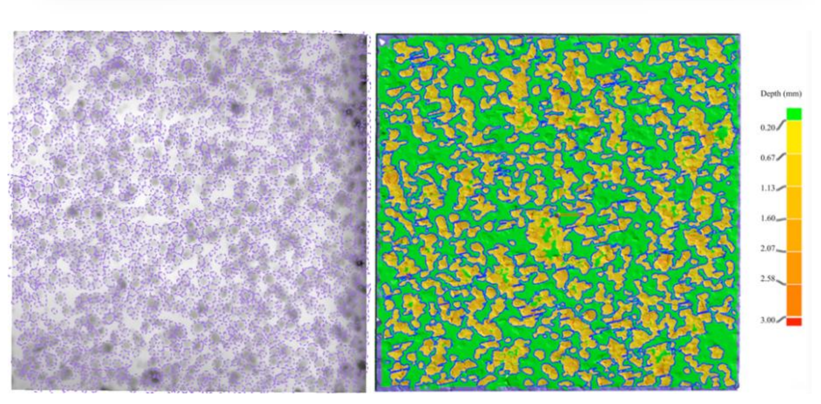

Automated U-Net model pipeline analysis of a densely-dented (saturated) hailpad (left) against automated, general-purpose third-party analysis (right).

Methods

The U-Net architecture is adapted here for instance segmentation through point-based querying to isolate individual dents. The hailpad (**1**) is first laser-scanned to produce a depth map (**2**), which is binarized using **3**–**5** (manually prescribed). A point query selection algorithm (**6**) identifies pixel locations for segmentation, after which a U-Net model processes the binary mask and point inputs (**7**). Erroneous predictions are removed using DBSCAN clustering. For each filtered dent, a minimal-area ellipse is fitted, and the maximum depth is retrieved from the darkest pixel in the depth map. Pixel measurements are converted to millimetres using the hailpad's dimensions and maximum surface depth, such that dent measurements are expressed in physical units.

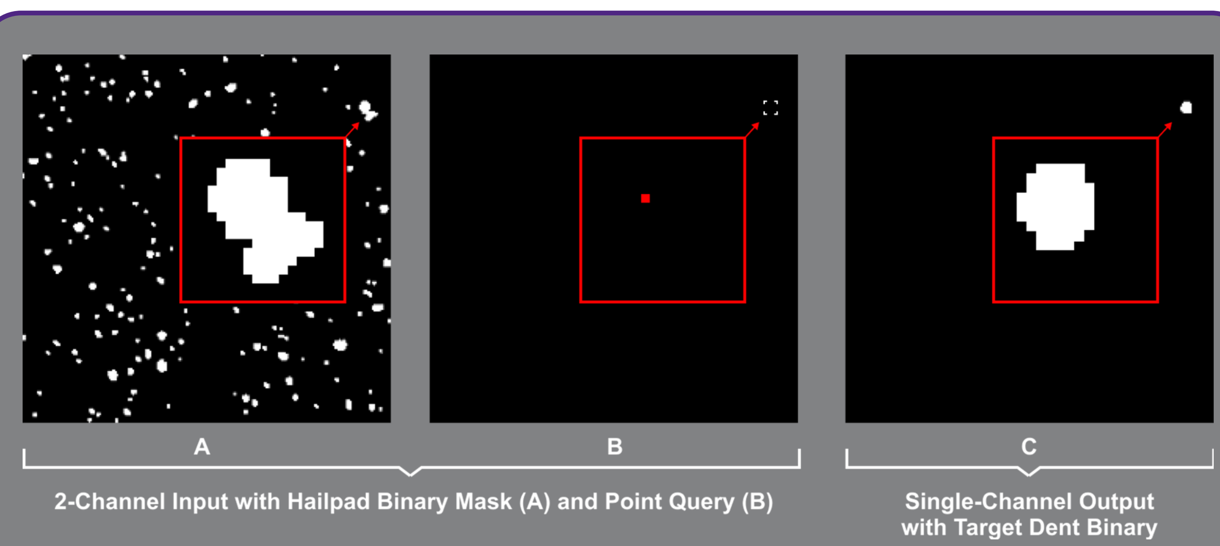

Input and output structure for U-net instance segmentation model. The input image consists of 2 channels, where the first channel (A) is the synthetic hailpad binary mask, and the second channel (B) is a single pixel value representing the point query. In the output image (C), only the target dent binary mask corresponding to the point query is shown.

Results

- With 1,635,000 synthetic point-mask samples (90% training; 10% validation), the U-Net model achieved a validation accuracy (IoU) score of 93.70% and 0.9652 Dice on predicted dents.
- Applied to 16 Alberta hailpads (2022–2024), it detected more individual dents by separating clustered impacts.
- Third-party methods (laboratory that specializes in high-precision, general-purpose 3D object scanning) merged clusters into fewer dents, while the U-Net gave finer detail but sometimes overestimated depth.
- Both methods struggle with small or shallow dents lost during thresholding or downsampling.

References

[1] Lozowski, E. P., and G. S. Strong, 1978b: On the calibration of hailpads. *J. Appl. Meteor.*, **17**, 521–528, https://doi.org/10.1175/1520-0450(1978)017<0521:OTCOH>2.0.CO;2.

[2] Zuiderveld, K., 1994: Contrast limited adaptive histogram equalization. In *Graphics Gems IV*, P. S. Heckbert, Ed., Academic Press, pp. 474–485, https://www.cse.unr.edu/~bebis/CS474/StudentPaperPresentations/1994%20-%20CLAHE.pdf.