Building Machine-Learning-Ready Data Sets for Martian Frost Detection

Gary Doran[1], Serina Diniega[1], Kiri L. Wagstaff[1], Mark Wronkiewicz[1], Steven Lu[1], and Jake Widmer [2]
[1]Jet Propulsion Laboratory, California Institute of Technology (Gary.B.Doran.Jr@jpl.nasa.gov),
[2]University of California Los Angeles

Introduction: Seasonal frost deposition and sublimation on the Martian surface is a key driver of Mars' atmospheric dynamics and is hypothesized to play a role in the evolution of geomorphological features on the surface. To study the impact of frost on these global and local processes, we seek to use machine learning (ML) to build frost detection models that can help scale up past manual studies in northern mid-latitudes in order to produce global monthly frost maps at a higher resolution (10-1000 meters). To improve the robustness and coverage of the frost map, we will combine the outputs of detection models applied towards instrument data observing in the visible, thermal, and spectral domains. Here, we describe lessons learned from developing a training set with ~1000 image tiles for a visible frost detection model using HiRISE [1] and CTX [2] images.

Translating Science Methodology: To translate this frost detection problem into one appropriate for ML, we need to capture a set of training data (labeled examples of frosted and unfrosted terrain). Preliminary efforts focused on previously studied northern mid-latitude (NML) regions where visible cues across multiple scales were used to identify frost. However, visible indicators can be ambiguous (e.g., a bright surface could indicate dust or frost) and robust interpretation often requires correlative analysis across indicators.. Thus, the first challenge in translating this nuanced science methodology into a simplified version that can be modeled with a binary classifier is determining a consistent and agreed upon set of criteria for identifying frost. We used three strategies to overcome this challenge:

1) Iterative Labeling Guide Development: First, we established an iterative process to define the manual labeling approach, which entails tracing polygonal annotations on top of images in the Labelbox.com platform (Figure 1). During early iterations, we met with domain experts and collaboratively walked through labeling example NML regions. After several iterations, data scientists on the project team applied what they had learned and performed practice 30-minute individual labeling sessions. The group met again with experts after individual sessions to compare labels and resolve ambiguities. During each iteration, clarifications to the labeling strategy were captured in a labeling guide, a product released with the training set that describes the labeling methodology and criteria in detail.

2) Contextualizing the Labeling Task: The second strategy used to simplify the labeling task is to leverage expertise from previous studies for which a high-level assessment of frost presence in each full observation has already been completed. This process incorporates contextual information used by the manual methodology for frost detection, and simplifies the labeling task conditioned on this *a priori* knowledge. Future work will investigate biases induced by sampling positive examples from these well-studied regions and negative examples based on physical contraints (latitude and season).

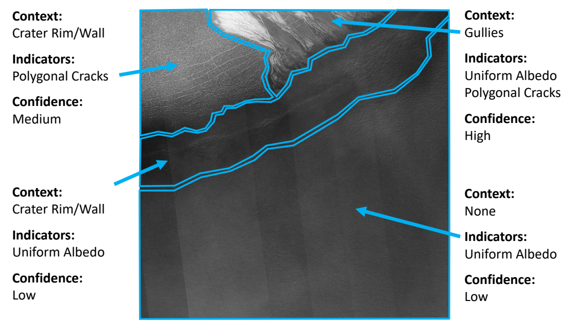

Figure 1: Frost annotation polygons labeled on top of a HiRISE image from Labelbox. Contextual information captured from the labeler is also shown.

3) Capturing Geologic Context: In addition to using Labelbox for determining the presence of frost indicators with polygonal regions, we also capture the confidence level and geologic context for each polygon including crater rim/wall, gullies, or dunes. From iterative labeling sessions, we learned to expect the detection task to be easier in certain contexts. We will use these annotations to investigate any systematic errors and biases on sub-categories of the data in trained ML models.

Conclusion: A goal of our task is to share not only technical results, but also our methodological experiences and lessons learned for other collaborations between scientists and data scientists. This abtract is intended to provide an example of how an ML-ready training set was constructed for Martian frost detection and provide insights for similar future collaborations.

References: [1] McEwen, A. (2007) *Mars Reconnaissance Orbiter High Resolution Imaging Science Experiment, Reduced Data Record*, 10.17189/1520303. [2] Malin, M. (2007) *MRO Context Camera Experiment Data Record Level 0 V1.0*, 10.17189/1520266.

Acknowledgments: Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.